

---

# Porting Applications from HP OpenVMS Alpha to HP OpenVMS Industry Standard 64 for Integrity Servers

Order Number: BA442–90001

**January 2005**

This manual provides a framework for application developers who are migrating applications from HP OpenVMS Alpha to HP OpenVMS Industry Standard 64 for Integrity Servers.

<b>Revision/Update Information:</b>	This is a new manual.
<b>Software Version:</b>	OpenVMS I64 Version 8.2
	OpenVMS Alpha Version 8.2

**Hewlett-Packard Company**  
**Palo Alto, California**

---

© Copyright 2005 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel, Itanium, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java is a U.S. trademark of Sun Microsystems, Inc.

UNIX is a registered trademark of The Open Group.

Printed in the US

ZK6673

This document was prepared using DECdocument, Version 3.3-1b.

---

# Contents

<b>Preface</b> .....	ix
<b>1 Introduction</b>	
1.1 OpenVMS Industry Standard 64 for Integrity Servers .....	1-1
1.2 Overview of the Porting Process .....	1-1
1.2.1 Evaluating the Application .....	1-1
<b>2 Fundamental Differences to Consider</b>	
2.1 Calling Standard .....	2-1
2.1.1 Changes to the OpenVMS Calling Standard .....	2-1
2.1.2 Extensions to the OpenVMS Calling Standard .....	2-3
2.2 Floating-Point Arithmetic .....	2-4
2.2.1 Overview .....	2-4
2.2.2 Impact on OpenVMS Applications .....	2-5
2.3 Object File Format .....	2-5
<b>3 Assessing Needs for Porting Your Application</b>	
3.1 Overview .....	3-1
3.2 Evaluating the Application .....	3-2
3.2.1 Selecting the Porting Method .....	3-2
3.2.2 Identifying Applications to Be Ported .....	3-4
3.2.3 Assessing Dependencies .....	3-5
3.2.3.1 Software Dependencies .....	3-5
3.2.3.2 Development Environment .....	3-6
3.2.3.3 Operating Environment .....	3-6
3.2.4 Operational Tasks .....	3-6
3.3 HP Porting Resources .....	3-8
3.4 Additional Considerations .....	3-8
<b>4 Migrating Source Modules</b>	
4.1 Setting Up the Migration Environment .....	4-2
4.1.1 Hardware .....	4-2
4.1.2 Software .....	4-3
4.1.2.1 HP OpenVMS Migration Software for Alpha to Integrity Servers and Translated Image Environment (TIE) .....	4-4
4.1.3 Coexistence with Translated Images .....	4-5
4.2 Compiling Applications on Alpha With Current Compiler Version .....	4-6
4.3 Testing Applications on Alpha for Baseline Information .....	4-7
4.4 Recompiling and Relinking on an I64 System .....	4-7
4.5 Debugging the Migrated Application .....	4-7

4.5.1	Debugging .....	4-7
4.5.2	Analyzing System Crashes .....	4-7
4.5.2.1	System Dump Analyzer .....	4-8
4.5.2.2	Crash Log Utility Extractor .....	4-8
4.6	Testing the Migrated Application .....	4-8
4.6.1	Alpha Tests Ported to I64 .....	4-8
4.6.2	New I64 Tests .....	4-9
4.6.3	Uncovering Latent Bugs .....	4-9
4.7	Integrating the Migrated Application into a Software System .....	4-9
4.8	Modifying Certain Types of Code .....	4-9
4.8.1	Conditionalized Code .....	4-10
4.8.1.1	MACRO Sources .....	4-10
4.8.1.2	BLISS Sources .....	4-10
4.8.1.3	C Sources .....	4-11
4.8.1.4	Existing Conditionalized Code .....	4-11
4.8.2	System Services With Alpha Architecture Dependencies .....	4-12
4.8.2.1	SYS\$GOTO_UNWIND .....	4-12
4.8.2.2	SYS\$LKWSET and SYS\$LKWSET_64 .....	4-12
4.8.3	Code With Other Dependencies on the Alpha Architecture .....	4-12
4.8.3.1	Initialized Overlaid Program Sections .....	4-12
4.8.3.2	Condition Handlers Use of SS\$_HPARITH .....	4-13
4.8.3.3	Mechanism Array Data Structure .....	4-13
4.8.3.4	Reliance on Alpha Object File Format .....	4-13
4.8.4	Code that Uses Floating-Point Data Types .....	4-13
4.8.4.1	LIB\$WAIT Problem and Solution .....	4-14
4.8.5	Incorrect Command Table Declaration .....	4-15
4.8.6	Code that Uses Threads .....	4-16
4.8.6.1	Thread Routines cma_delay and cma_time_get_expiration .....	4-17
4.8.7	Code With Unaligned Data .....	4-17
4.8.8	Code that Relies on the OpenVMS Alpha Calling Standard .....	4-18
4.8.9	Privileged Code .....	4-19
4.8.9.1	Use of SYS\$LKWSET and SYS\$LKWSET_64 .....	4-19
4.8.9.2	Use of SYS\$LCKPAG and SYS\$LCKPAG_64 .....	4-20
4.8.9.3	Terminal Drivers .....	4-20
4.8.9.4	Protected Image Sections .....	4-20

## 5 OpenVMS I64 Development Environment

5.1	Native I64 Compilers .....	5-1
5.1.1	VAX MACRO-32 Compiler for OpenVMS I64 .....	5-1
5.2	Other Development Tools .....	5-2
5.2.1	Translating Alpha Code .....	5-3
5.3	Linking Modules .....	5-3
5.3.1	Differences When Linking on OpenVMS I64 Systems .....	5-4
5.3.1.1	No Based Clusters .....	5-4
5.3.1.2	Handling of Initialized Overlaid Program Sections on OpenVMS I64 .....	5-5
5.3.1.3	Behavior Difference When Linking ELF Common Symbols .....	5-5
5.3.2	Expanded Map File Information .....	5-6



5.3.3	New Linker Qualifiers and Options for OpenVMS I64 .....	5-6
5.3.3.1	New /BASE_ADDRESS Qualifier .....	5-6
5.3.3.2	New /SEGMENT_ATTRIBUTE Qualifier .....	5-6
5.3.3.3	New /FP_MODE Qualifier .....	5-7
5.3.3.4	New /EXPORT_SYMBOL_VECTOR and /PUBLISH_GLOBAL_SYMBOLS Qualifiers .....	5-7
5.3.3.5	New Alignments for the PSECT_ATTRIBUTE Option .....	5-8
5.3.3.6	New GROUP_SECTIONS and SECTION_DETAILS keywords for the /FULL Qualifier .....	5-8
5.3.4	Mixed-Case Arguments in Linker Options, Revisited .....	5-8
5.4	Debugging Capabilities on OpenVMS I64 Systems .....	5-9
5.4.1	OpenVMS Debugger .....	5-9
5.4.1.1	Architecture Support .....	5-9
5.4.1.2	Language Support .....	5-10
5.4.1.3	Functional Areas and Commands .....	5-10
5.4.1.4	Functionality Not Yet Ported .....	5-11
5.4.2	XDelta Debugger .....	5-11
5.4.2.1	XDelta Capabilities on OpenVMS I64 .....	5-11
5.4.2.2	Differences Between XDelta on OpenVMS I64 and OpenVMS Alpha Systems .....	5-11
5.5	OpenVMS I64 Librarian Utility .....	5-12
5.5.1	Considerations When Using the I64 Librarian .....	5-12
5.5.2	Changes to the LBR\$ Routines .....	5-12
5.5.3	I64 Library Format Handles UNIX-Style Weak Symbols .....	5-13
5.5.3.1	New ELF Type for Weak Symbols .....	5-13
5.5.3.2	Version 6.0 Library Index Format .....	5-13
5.5.3.3	New Group-Section Symbols .....	5-13
5.5.3.4	Current Library Limitation with Regard to Weak and Group Symbols .....	5-13

## 6 Preparing to Port Applications

6.1	Ada .....	6-2
6.2	BASIC .....	6-2
6.3	BLISS Compiler .....	6-2
6.3.1	BLISS File Types and File Location Defaults .....	6-2
6.3.2	Alpha BLISS Features Not Available .....	6-3
6.3.3	OpenVMS I64 BLISS Features .....	6-4
6.4	COBOL .....	6-15
6.4.1	Floating-Point Arithmetic .....	6-15
6.4.2	/ARCH and /OPTIMIZE=TUNE Qualifiers .....	6-15
6.5	Fortran .....	6-15
6.5.1	Floating-Point Arithmetic .....	6-15
6.5.2	Only the F90 Compiler is Supported .....	6-17
6.5.3	/ARCH and /OPTIMIZE=TUNE Qualifiers .....	6-18
6.6	HP C/ANSI C .....	6-18
6.6.1	I64 Floating-Point Default .....	6-18
6.6.2	Semantics of /IEEE_MODE Qualifier .....	6-18
6.6.3	Predefined Macros .....	6-19
6.7	HP C++ .....	6-19
6.7.1	Floating Point and Predefined Macros .....	6-19
6.7.2	Long Double .....	6-20
6.7.3	Object Model .....	6-20
6.7.4	Language Dialects .....	6-20

6.7.5	Templates .....	6-20
6.8	Java .....	6-20
6.9	Macro-32 .....	6-20
6.10	HP Pascal .....	6-21
6.10.1	/ARCH and /OPTIMIZE=TUNE Qualifiers .....	6-21

## 7 Other Considerations

7.1	Hardware Considerations .....	7-1
7.1.1	Intel Itanium Processor Family Overview .....	7-1
7.1.2	Alpha Processor Family Overview .....	7-1
7.2	Endianism Considerations .....	7-2

## A Application Evaluation Checklist

## B Unsupported Layered Products

## C Porting Application-Specific Stack-Switching Code to I64

C.1	Overview of KP Services .....	C-1
C.1.1	Terminology .....	C-2
C.1.2	Stacks and Data Structures .....	C-2
C.1.3	KPBs .....	C-3
C.1.4	Supplied KPB Allocation Routines .....	C-4
C.1.5	Kernel Mode Allocation .....	C-4
C.1.6	Mode-Independent Allocation .....	C-6
C.1.7	Stack Allocation APIs .....	C-8
C.1.8	System-Supplied Allocation and Deallocation Routines .....	C-8
C.1.9	End Routine .....	C-9
C.2	KP Control Routines .....	C-9
C.2.1	Overview .....	C-9
C.2.2	Routine Descriptions .....	C-11
C.2.2.1	EXE\$KP_START .....	C-11
C.2.2.2	EXE\$KP_STALL_GENERAL .....	C-11
C.2.2.3	EXE\$KP_RESTART .....	C-11
C.2.2.4	EXE\$KP_RESTART .....	C-12
C.3	Design Considerations .....	C-12

## Glossary

## Index

## Figures

3-1	Porting an Application .....	3-3
4-1	Migrating an Application from Alpha to I64 .....	4-1
7-1	Big-Endian Byte Ordering .....	7-3
7-2	Little-Endian Byte Ordering .....	7-3
C-1	KP Routine Execution .....	C-10

## Tables

2-1	OpenVMS Calling Standard Changes .....	2-1
2-2	OpenVMS Calling Standard Extensions .....	2-4
4-1	Compiler Switches for Reporting Compile-Time Reference .....	4-18
5-1	OpenVMS Development Tools .....	5-2
6-1	Mapping Alpha Compiler Versions to I64 Compiler Versions .....	6-1
B-1	Layered Products Not Ported to OpenVMS I64 .....	B-1
C-1	Allocation Guidelines by Mode and Scope .....	C-3
C-2	Kernel Mode Allocation Flags .....	C-5
C-3	Mode-Independent Allocation Flags .....	C-6





---

# Preface

## Intended Audience

*Porting Applications from HP OpenVMS Alpha to HP OpenVMS Industry Standard 64 for Integrity Servers* is intended for application developers who are planning to migrate applications from the OpenVMS Alpha operating system to the HP OpenVMS Industry Standard 64 for Integrity Servers (I64) operating system.

## Document Structure

This document is organized as follows:

Chapter 1 includes an overview of the OpenVMS I64 8.2 operating system, including its evolution and its benefits.

Chapter 2 discusses fundamental differences between the OpenVMS Alpha and OpenVMS I64 operating systems.

Chapter 3 outlines a process for assessing applications in preparation for porting to OpenVMS I64.

Chapter 4 discusses specific tasks that must be completed to prepare source modules for migration, as well as information about compiling, linking, testing, and deploying ported applications.

Chapter 5 discusses the OpenVMS I64 development environment.

Chapter 6 discusses porting considerations related to the primary compilers that will be available for OpenVMS I64.

Chapter 7 discusses other factors to consider when undertaking a porting effort.

Appendix A contains a sample checklist for use when evaluating applications prior to porting.

Appendix B lists the HP layered products that are supported on OpenVMS Alpha but not on OpenVMS I64.

Appendix C describes how to use KP routines when porting application-specific stack-switching code to I64.

The Glossary defines important terms and abbreviations used in this manual.

## Related Documents

Users of this manual may also find the following documents useful:

- *HP OpenVMS Calling Standard*
- *HP OpenVMS Debugger Manual*
- *HP OpenVMS Linker Utility Manual*

- *HP OpenVMS System Analysis Tools Manual*

Documentation for individual compilers may also be useful in the porting process.

The “OpenVMS Floating-Point Arithmetic on the Intel® Itanium® Architecture” white paper describes differences between floating-point data types on OpenVMS Alpha and OpenVMS I64. You can obtain the white paper from the following location:

[http://www.hp.com/products1/evolution/alpha\\_retaintrust/openvms/resources.html](http://www.hp.com/products1/evolution/alpha_retaintrust/openvms/resources.html)

For additional information about HP OpenVMS products and services, visit the following World Wide Web address:

<http://www.hp.com/go/openvms>

## Reader's Comments

HP welcomes your comments on this manual. Please send comments to either of the following addresses:

Internet	<b>openvmsdoc@hp.com</b>
Postal Mail	Hewlett-Packard Company OSSG Documentation Group, ZKO3-4/U08 110 Spit Brook Rd. Nashua, NH 03062-2698

## How to Order Additional Documentation

For information about how to order additional documentation, visit the following World Wide Web address:

<http://www.hp.com/go/openvms/doc/order>

## Conventions

The following conventions may be used in this manual:

Ctrl/ <i>x</i>	A sequence such as Ctrl/ <i>x</i> indicates that you must hold down the key labeled Ctrl while you press another key or a pointing device button.
PF1 <i>x</i>	A sequence such as PF1 <i>x</i> indicates that you must first press and release the key labeled PF1 and then press and release another key or a pointing device button.
<div style="border: 1px solid black; padding: 2px;">Return</div>	<p>In examples, a key name enclosed in a box indicates that you press a key on the keyboard. (In text, a key name is not enclosed in a box.)</p> <p>In the HTML version of this document, this convention appears as brackets, rather than a box.</p>

...	<p>A horizontal ellipsis in examples indicates one of the following possibilities:</p> <ul style="list-style-type: none"> <li>• Additional optional arguments in a statement have been omitted.</li> <li>• The preceding item or items can be repeated one or more times.</li> <li>• Additional parameters, values, or other information can be entered.</li> </ul>
.	A vertical ellipsis indicates the omission of items from a code example or command format; the items are omitted because they are not important to the topic being discussed.
( )	In command format descriptions, parentheses indicate that you must enclose choices in parentheses if you specify more than one.
[ ]	In command format descriptions, brackets indicate optional choices. You can choose one or more items or no items. Do not type the brackets on the command line. However, you must include the brackets in the syntax for OpenVMS directory specifications and for a substring specification in an assignment statement.
	In command format descriptions, vertical bars separate choices within brackets or braces. Within brackets, the choices are optional; within braces, at least one choice is required. Do not type the vertical bars on the command line.
{ }	In command format descriptions, braces indicate required choices; you must choose at least one of the items listed. Do not type the braces on the command line.
<b>bold type</b>	Bold type represents the introduction of a new term. It also represents the name of an argument, an attribute, or a reason.
<i>italic type</i>	Italic type indicates important information, complete titles of manuals, or variables. Variables include information that varies in system output (Internal error <i>number</i> ), in command lines (/PRODUCER= <i>name</i> ), and in command parameters in text (where <i>dd</i> represents the predefined code for the device type).
UPPERCASE TYPE	Uppercase type indicates a command, the name of a routine, the name of a file, or the abbreviation for a system privilege.
Example	This typeface indicates code examples, command examples, and interactive screen displays. In text, this type also identifies URLs, UNIX commands and pathnames, PC-based commands and folders, and certain elements of the C programming language.
-	A hyphen at the end of a command format description, command line, or code line indicates that the command or statement continues on the following line.
numbers	All numbers in text are assumed to be decimal unless otherwise noted. Nondecimal radices—binary, octal, or hexadecimal—are explicitly indicated.



---

# Introduction

This chapter contains an introduction to OpenVMS Industry Standard 64 for Integrity Servers (I64) systems and an overview of the migration process.

For more information, refer to the *HP OpenVMS Version 8.2 Release Notes*.

## 1.1 OpenVMS Industry Standard 64 for Integrity Servers

The OpenVMS I64 architecture has a 64-bit model and basic system functions similar to the Alpha architecture; thus the great majority of applications running on OpenVMS Alpha today can be easily ported to the I64 architecture with few changes.

The OpenVMS Alpha and OpenVMS I64 architecture variants of OpenVMS are produced from a single-source code base; thus, nonhardware dependent features may be incorporated into both versions without multiple changes to the source code, thereby minimizing the time required to perform qualification testing, and helping to ensure the availability of critical applications on both OpenVMS platforms.

HP intends to maintain a strict policy of ensuring forward source compatibility so that “well-behaved” applications that currently run on recommended versions of OpenVMS Alpha run successfully on OpenVMS I64 systems. If an application takes advantage of published system services and library interfaces, there should be a high level of confidence that the application will move without modifications to the latest version of OpenVMS I64.

OpenVMS I64 has the same look and feel familiar to OpenVMS customers. Minor changes were made to accommodate the new architecture, but the basic structure and capabilities of OpenVMS are the same.

## 1.2 Overview of the Porting Process

This section provides an overview of the porting process as it applies to most applications.

### 1.2.1 Evaluating the Application

Evaluating the application identifies the steps necessary for its migration to HP OpenVMS I64. The result of the evaluation should be a migration plan that answers the following questions:

- How to migrate the application
- How much time, effort, and cost the migration requires
- Whether you require support from HP services

Evaluation should include the following three stages:

1. A general inventory of the application, including identifying dependencies on other software

## Introduction

### 1.2 Overview of the Porting Process

2. Source analysis to identify dependencies on other software
3. Selection of a migration method (rebuilding from sources or using a binary translator)

Completing these steps yields the information necessary to write an effective migration plan.

The first step in evaluating an application for migration is to determine exactly what has to be migrated. This includes not only the application itself, but everything that the application requires in order to run properly. To begin evaluating your application, identify and locate the following items:

- Parts of the application
  - Source modules for the main program
  - Shareable images
  - Object modules
  - Libraries (object module, shareable image, text, or macro)
  - Data files and databases
  - Message files
  - Documentation
- Other software on which your application depends; such as:
  - Run-time libraries
  - HP layered products
  - Third-party layered products
- Required operating environment
  - System characteristics

What sort of system is required to run and maintain your application? For example, how much memory is required, how much disk space, and so on?
  - Build procedures

This includes HP tools such as Code Management System (CMS) and Module Management System (MMS).
  - Test Suite

Your tests must be able both to confirm that the migrated application runs correctly and to evaluate its performance.

Many of these items have already been migrated to OpenVMS I64; for example:

- HP software bundled with the OpenVMS operating system
  - Run-time libraries
  - Other shareable libraries, such as those supplying callable utility routines and other application library routines
- HP layered products
  - Compilers and compiler RTLs
  - Database managers



## 1.2 Overview of the Porting Process

- Networking environment
- Third-party products

Many third-party products now run on OpenVMS I64. To determine whether a particular application has been migrated, contact the application vendor.

You are responsible for migrating your application and your development environment, including build procedures and test suites.

When you have completed the evaluation of your application, continue with a more detailed evaluation of each module and image, as described in Chapter 4.



## Fundamental Differences to Consider

This chapter discusses the fundamental differences between the Alpha and I64 architectures.

### 2.1 Calling Standard

As with other components, the implementation of the OpenVMS Calling Standard on the Intel® Itanium® processor family sought to differ as little as possible from the OpenVMS VAX and Alpha conventions. The design methodology started with Itanium conventions and made changes only where necessary. This helped to maintain compatibility with historical OpenVMS design while minimizing the cost and difficulty of porting applications and the operating system itself to the Itanium architecture.

The following sections contain a high-level description of differences between the *Itanium® Software Conventions and Runtime Architecture Guide* and the OpenVMS Calling Standard. For more detailed information, see the *HP OpenVMS Calling Standard*.

#### 2.1.1 Changes to the OpenVMS Calling Standard

Table 2–1 describes the major changes to the OpenVMS Calling Standard to for OpenVMS I64.

**Table 2–1 OpenVMS Calling Standard Changes**

Item	Description
Data model	OpenVMS on Alpha systems is deliberately ambiguous about the data model in use: many programs are compiled using what appears to be an ILP32 model, yet most of the system operates as though using either a P64 or LP64 model. The sign extension rules for integer parameters play a key role in making this more or less transparent. OpenVMS I64 preserves this characteristic, while the Itanium architecture conventions define a pure LP64 data model.
Data terminology	This specification uses the terms <i>word</i> and <i>quadword</i> to mean 2 bytes and 8 bytes, respectively, while the Itanium terminology uses these words to mean 4 bytes and 16 bytes, respectively.

(continued on next page)

## Fundamental Differences to Consider

### 2.1 Calling Standard

**Table 2–1 (Cont.) OpenVMS Calling Standard Changes**

Item	Description
General register usage	<p>General registers are used for integer arithmetic, some parts of VAX floating-point emulation, and other general-purpose computation. OpenVMS uses the same (default) conventions for these registers except for the following cases:</p> <ul style="list-style-type: none"><li>• R8 and R9 (only) are used for return values.</li><li>• R10 and R11 are used as scratch registers and not for return values.</li><li>• R25 is used for an AI (argument information) register.</li></ul>
Floating-point register usage	<p>Floating-point registers are used for floating-point computations, some parts of VAX floating-point emulation, and certain integer computations. OpenVMS uses the same (default) conventions for these registers except in the following cases:</p> <ul style="list-style-type: none"><li>• F8 and F9 (only) are used for return values.</li><li>• F10 through F15 are used as scratch registers and not for return values.</li></ul>

(continued on next page)

**Table 2–1 (Cont.) OpenVMS Calling Standard Changes**

Item	Description
Parameter passing	<p>OpenVMS parameter passing is similar to the Itanium conventions. Note the following differences in the OpenVMS standard:</p> <ul style="list-style-type: none"> <li>• The OpenVMS standards adds an argument information register (for argument count and parameter type information).</li> <li>• No argument is ever duplicated in both general and floating-point registers.</li> <li>• For parameters that are passed in registers, the first parameter is passed in either the first general-register slot (R32) or the first floating-point register slot (F16), the second parameter in either the second general register (R33) or second floating-point register (F17) slot, and so on. Floating-point parameters are not packed into the available floating-point registers, and a maximum of eight parameters total are passed in registers.</li> <li>• For 32-bit parameters passed in the general registers, the 32-bit value is sign extended to the full 64-bit width of the parameter slot by replicating bit 31 (even for unsigned types).</li> <li>• There is no even slot alignment for arguments larger than 64 bits.</li> <li>• There is no special handling for HFA (homogeneous floating-point aggregates) in general, although some rules for complex types have a similar benefit.</li> <li>• OpenVMS implements <code>__float128</code> pass-by value semantics using a reference mechanism.</li> <li>• OpenVMS supports only little-endian representations.</li> <li>• OpenVMS supports three additional VAX floating-point types for backward compatibility: <code>F_floating</code> (32 bits), <code>D_floating</code> (64 bits), and <code>G_floating</code> (64 bits). Values of these types are passed using the general registers.</li> </ul>
Return values	<p>Return values up to at most 16 bytes in size may be returned in registers; larger return values are returned using a <b>hidden parameter</b> method using the first or second parameter slot.</p>

### 2.1.2 Extensions to the OpenVMS Calling Standard

Table 2–2 describes additions to or extensions of the OpenVMS Calling Standard.

## Fundamental Differences to Consider

### 2.1 Calling Standard

**Table 2–2 OpenVMS Calling Standard Extensions**

Item	Description
Floating-point data types	The OpenVMS calling standard includes support for the VAX <code>F_floating</code> (32-bit), <code>D_floating</code> (64-bit), and <code>G_floating</code> (64-bit) data types found on VAX and Alpha systems; it omits support for the 80-bit double-extended floating-point type of the Itanium architecture.
VAX compatible record layout	The OpenVMS standard adds a user-optional, VAX-compatible record layout.
Linkage options	OpenVMS allows additional flexibility and user control in the use of the static general registers as inputs, outputs, global registers and whether used at all.
Memory stack overflow checking	OpenVMS defines how memory stack overflow checking should be performed.
Function descriptors	OpenVMS defines extended forms of function descriptors to support additional functionality for bound procedure values and translated image support.
Unwind information	OpenVMS adds an operating system-specific data area to the unwind information block of the Itanium architecture. The presence of an operating system-specific data area is indicated by a flag in the unwind information header.
Handler invocation	OpenVMS does not invoke a handler while control is in either a prologue or epilogue region of a routine. This difference in behavior is indicated by a flag in the unwind information header.
Translated images	OpenVMS adds support (signature information and special ABIs) for calls between native and translated VAX or Alpha images.

## 2.2 Floating-Point Arithmetic

This section discusses the differences in floating-point arithmetic on OpenVMS VAX, OpenVMS Alpha, and OpenVMS I64 systems.

### 2.2.1 Overview

The Alpha architecture supports both IEEE and VAX floating-point formats in hardware. OpenVMS compilers generate code using the VAX formats by default, with options on Alpha to use IEEE formats. The Itanium architecture implements floating-point arithmetic in hardware using the IEEE floating-point formats, including IEEE single and IEEE double.

If an application was originally written for OpenVMS VAX or OpenVMS Alpha using the default floating-point formats, it can be ported to OpenVMS I64 in one of two ways: continue to use VAX floating-point formats utilizing the conversion features of the HP compilers or convert the application to use IEEE floating-point formats. VAX floating-point formats can be used in those situations where access to previously generated binary floating-point data is required. HP compilers generate the code necessary to convert the data between VAX and IEEE formats.

For details about the conversion process, see the “OpenVMS Floating-Point Arithmetic on the Intel® Itanium® Architecture” white paper. See the Related Documents section in the Preface for the web location of this white paper.



IEEE floating-point format should be used in those situations where VAX floating-point formats are not required. The use of IEEE floating-point formats will result in more efficient code.

### **2.2.2 Impact on OpenVMS Applications**

VAX floating-point formats are supported on the Itanium architecture by converting them to IEEE single and IEEE double floating types. By default, this is a transparent process that will not impact most applications. HP is providing compilers for C, C++, Fortran, BASIC, Pascal, and COBOL, all with the same floating-point options. Application developers need only recompile their applications using the appropriate compiler. For detailed descriptions of floating-point options, refer to the individual compilers' documentation.

Because IEEE floating-point format is the default, unless an application explicitly specifies VAX floating-point format options, a simple rebuild for OpenVMS I64 uses native IEEE formats directly. For programs that do not depend directly on the VAX formats for correct operation, this is the most desirable way to build for OpenVMS I64.

## **2.3 Object File Format**

Hewlett-Packard has determined that the most efficient way to utilize the compiler technologies developed by Intel is to adopt two industry standards, both of which are used by the Itanium compiler:

- ELF (executable and linkable format) describes the object file and executable file formats.
- DWARF (debugging and traceback information format) describes the way debugging and traceback information is stored in the object and executable files.

One clear advantage of this decision is that development tools can be ported to OpenVMS more easily in the future. All of the compilers, development tools, and utilities provided with the OpenVMS I64 operating system will utilize and understand these new formats. Application developers will not need to be concerned with these formats unless their applications have specific knowledge of them.

In a typical application scenario in which an application is compiled, linked, debugged, and deployed, the details of the object file format, the executable image file format, and the debugging and traceback information are of no concern to the developer. However, software engineers who develop compilers, debuggers, analysis tools, or other utilities that are dependent on these file formats will need to understand how they are being implemented on OpenVMS.



---

## Assessing Needs for Porting Your Application

Typically, porting an application to a new platform involves the following procedure:

1. Assessing porting needs
2. Identifying architecture dependencies and nonstandard coding practices in the application to be ported, and rewriting code accordingly
3. Compiling, linking, and running the application
4. Performing testing before and after porting
5. Recompiling the code and repeating steps 1 through 4, as necessary
6. Shipping the revised application
7. Operating and installing processes or procedures
8. Supporting and maintaining the application

Although you complete these tasks sequentially, you might have to repeat some steps if problems arise.

This chapter addresses the first task — assessing porting needs. Subsequent chapters discuss the remaining tasks. This chapter describes the responsibilities and issues that developers should consider in planning to port an application. In particular, it sets the context for porting applications from an OpenVMS Alpha to an OpenVMS I64 environment and provides the information necessary to help developers evaluate needs, plan, and begin porting software.

Appendix A provides a checklist to guide you through the planning and porting process.

### 3.1 Overview

Most applications running on OpenVMS Alpha today can easily be ported to OpenVMS I64 with few changes. The Alpha and I64 variants of OpenVMS are produced from a single source-code base, enabling developers to incorporate features (that are independent of the hardware) into both versions without requiring multiple changes to the source code. This minimizes the time required to perform qualification testing and helps to ensure the availability of critical applications on both OpenVMS platforms.

HP intends to maintain a strict policy of ensuring and maintaining forward source compatibility so that “well-behaved” applications that run on recommended versions of OpenVMS Alpha will also run successfully on OpenVMS I64 systems. For the most part, if an application takes advantage of published system services and library interfaces, the application should be portable (as is) to the latest version of OpenVMS I64. However, some published system and library interfaces available on OpenVMS Alpha will not be available or will behave differently on OpenVMS I64. In addition, the Linker utility does not support all qualifiers

## Assessing Needs for Porting Your Application

### 3.1 Overview

supported on OpenVMS Alpha. For more information about these exceptions, see Chapter 4 and Chapter 5.

OpenVMS I64 has the same look and feel familiar to OpenVMS customers. Minor changes were needed to accommodate the new architecture, but the basic structure and capabilities of OpenVMS are the same.

### 3.2 Evaluating the Application

The first step to porting an application is to evaluate and identify the steps necessary to port the application to OpenVMS I64. The evaluation process culminates in a porting plan that answers the following questions:

- How do you port the application?
- How much time, effort, and cost does porting require?
- Do you require support from HP services?

The evaluation process has the following four steps:

1. Identifying applications to be ported
2. Identify dependencies on other software
3. Selecting a porting method
4. Analyzing operational tasks

Completing these steps yields the information necessary to write an effective porting plan. Sections 3.2.1 through 3.2.4 suggest the steps for an initial, rudimentary evaluation, after which you can perform a more detailed evaluation of each module, image, and all dependencies.

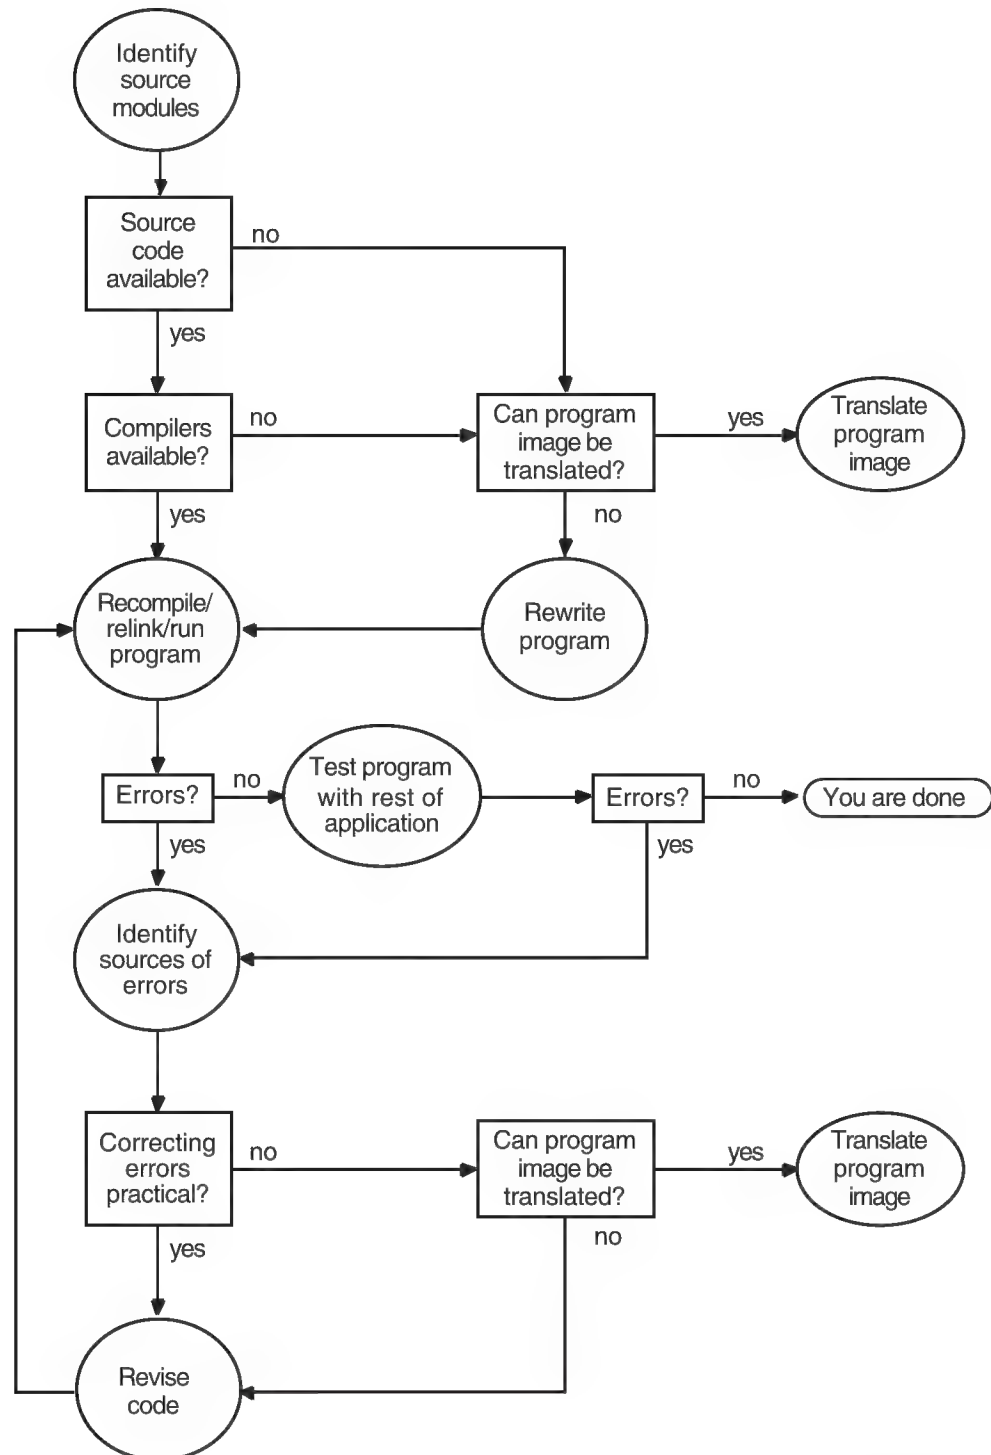
After you complete the evaluation steps described in the following sections, compare your results with the supported product, procedures, and functionality of the target platform. Research any deviations for schedule mismatches, missing functionality, procedural conflicts, and supportability. To complete your porting plan, include all costs and time impacts.

#### 3.2.1 Selecting the Porting Method

After you complete the evaluation process described in the following sections, you will have to port your application and your development environment. In addition to the source modules, you might have to port other components, such as build procedures, test suites, and in some cases, the application data structures. You must decide on the best method for porting each part of the application. Usually the issue is whether to rebuild your application from sources or to use a binary translator. To make this decision, you need to know which methods are possible for each part of the application, and how much work is required for each method. To help answer those questions, you should answer the series of questions and perform the tasks shown in Figure 3-1.

The majority of applications can be ported by recompiling and relinking them. If your application runs only in user mode and is written in a standard high-level language, it is most likely in this category.

Figure 3–1 Porting an Application



ZK-4990A-AI

## Assessing Needs for Porting Your Application

### 3.2 Evaluating the Application

#### 3.2.2 Identifying Applications to Be Ported

The first step in evaluating an application for porting is to identify exactly what has to be ported. This includes not only the application itself, but everything that the application requires to run properly. To begin evaluating your application, identify and locate the following items:

- Source modules for the main program
- Shareable images
- Object modules
- Libraries (object module, shareable image, text, or macro)
- Data files and databases
- Message files
- Scripts and procedures (build files)
- Application documentation (if any)

Consider the differences between OpenVMS Alpha and OpenVMS I64 described in Chapter 2 and the consequent changes that are necessary for application components prior to porting. In addition, consider the processor-related development issues, as described in Chapter 7.

Evaluate the time and costs required for these changes. Any code that depends specifically on the OpenVMS Alpha architecture must be changed. For example, an application may require changes if it includes code that:

- Deals with specific machine instructions or that makes assumptions about the number of registers or the functions of specific registers.
- Relies on particular elements of the OpenVMS Alpha calling standard. (For differences between OpenVMS Alpha and OpenVMS I64 Calling Standards, see Section 2.1.1.)
- Relies on unpublished system services and library interfaces, or certain system and library interfaces that are not available or will behave differently on OpenVMS I64. (For more information about the availability and divergent behavior of certain system and library interfaces, see Chapter 4 and Chapter 5.)
- Relies on the OpenVMS Alpha object file format, the executable image file format, or the debug symbol table format.
- Is conditionalized or includes logic that assumes it is running on either an OpenVMS VAX or an OpenVMS Alpha system which might not be able to handle a third platform.
- Depends on OpenVMS Alpha or VAX internal data structures; for example:
  - Floating-point formats that are not IEEE standard. (IEEE floating-point formats are the default for OpenVMS I64 and should be used where VAX floating-point formats are not required; for more information, see Section 3.2.4 and Section 4.8.4.)
  - The mechanism array data structure on OpenVMS I64 differs significantly from that on OpenVMS Alpha (see Section 4.8.3.3).



---

#### Note

---

HP recommends that you align your data naturally to achieve optimal performance for data referencing. On both OpenVMS Alpha and OpenVMS I64 systems, referencing data that is unaligned degrades performance significantly. In addition, unaligned shared data can cause a program to execute incorrectly. You must align shared data naturally. Shared data might be between threads of a single process, between a process and ASTs, or between several processes in a global section. For more information about data alignment, see Section 4.8.7 and the *HP OpenVMS Programming Concepts Manual*.

---

- References terminal drivers that do not use the call interface; for example, those using jump to subroutine (JSB) macros need to be revised to use the call interface
- Contains user-written threading-like, such as code that performs direct stack switching, that implements a co-routine or tasking model, or that depends on the nature of the procedure call stack frames. For more details, see Section 4.8.6.
- Incorporates nonstandard or undocumented code practices or interfaces.

In addition any source code that does not have a supported compiler on OpenVMS I64 must be rewritten or translated. (For a list of supported compilers, see Chapter 6.) Note also that certain language capabilities of OpenVMS Alpha are not supported on OpenVMS I64; for more information, see the *HP OpenVMS Version 8.2 Release Notes*. In addition, applications that run in privileged mode might require changes.

For more details about changes required in your source modules prior to porting them, see Chapter 4.

### 3.2.3 Assessing Dependencies

The next step is to assess the software and environment on which your application depends.

#### 3.2.3.1 Software Dependencies

Software dependencies might include such components as:

- Run-time libraries
- HP layered products
- Third-party layered products
- Tools

Many of these items have already been migrated to OpenVMS I64, including the following:

- HP software bundled with the OpenVMS operating system, including:
  - Run-time libraries
  - Other shareable libraries, such as those supplying callable utility routines and other application library routines
- HP layered products, including:
  - Compilers and compiler RTLs

## Assessing Needs for Porting Your Application

### 3.2 Evaluating the Application

- Database managers
- Networking environment
- Third-party products

Many third-party products now run on OpenVMS I64. To determine whether a particular application has been migrated, contact the application vendor.

---

#### Important

---

The availability of third-party software can be the biggest obstacle for a porting project. You must determine the availability and cost of third-party software for the target operating system. This might include build environment tools, third-party libraries, and automated testing tools. For information about tools that are freely available on OpenVMS, see Chapter 5.

---

#### 3.2.3.2 Development Environment

Consider the development environment upon which your application depends, including the operating system configuration, hardware, development utilities, build procedures (including HP tools such as Code Management System [CMS] and Module Management System [MMS]), and so forth.

#### 3.2.3.3 Operating Environment

Consider the following issues regarding the operating environment:

- System characteristics

What sort of system is required to run and maintain your application? For example, how much memory is required, how much disk space, and so on?
- Test suite

Your tests should confirm that the ported application runs correctly, and they should evaluate its performance. Regression testing is crucial to the development of any new product, especially to the porting of an application to a new system. Regression tests aid in testing software on new versions of the operating system and platform. The regression tests can be any one or combination of the following:

  - A software program written by your development team to exercise all possible code paths and functions
  - A series of DCL command procedures
  - Interactive testing conducted manually (instead of automatically by software)
  - HP Digital Test Manager or an equivalent Open Source tool

#### 3.2.4 Operational Tasks

Evaluate the responsibilities for and needs for operational tasks required to port and maintain the applications, for example:

- Installation requirements
- Compilation requirements
  - Is the compiler that you use for your application also supported on this version of OpenVMS I64?

## Assessing Needs for Porting Your Application

### 3.2 Evaluating the Application

- Prior to porting your application, HP recommends that you compile your application on an OpenVMS Alpha system using the latest version of the OpenVMS Alpha compiler. In this way, you can uncover problems that you might encounter when you compile the ported application on OpenVMS I64. Some newer versions of compilers apply a stricter interpretation of existing compiler standards or enforce newer, stricter standards. For more information about compilers supported on this version of OpenVMS I64, see Chapter 6.
- After testing with a compile on the OpenVMS Alpha system and resolving any problems, you need to port the application to your OpenVMS I64 system, recompile, relink, and requalify. Plan for these activities accordingly.
- Can you switch to IEEE floating point data types from VAX floating-point data types? In general, applications can still use VAX floating point-format on OpenVMS I64. Use VAX floating point formats when access to previously generated binary floating-point data is required. VAX floating point is achieved on OpenVMS I64 systems in software. HP compilers automatically generate the code necessary to convert VAX floating-point values to the IEEE floating-point format and perform the mathematical operation; however, they then convert the value back to VAX floating-point. This additional conversion adds execution overhead to the application. For this reason, consider switching your application to use IEEE floating-point format.

Applications on OpenVMS Alpha can be switched to use IEEE floating point with no performance penalty because IEEE floating-point is supported by the Alpha hardware. This will make migration of your application to OpenVMS I64 easier. IEEE floating-point is the only floating-point format supported by the Itanium hardware. In addition, your application will perform better.

You can test an application's behavior with IEEE floating-point values by compiling it on an OpenVMS Alpha system with an IEEE qualifier. (For many compilers, specify the compiler option `/FLOAT=IEEE`. If using I64 BASIC, use the `/REAL_SIZE` qualifier.) If that produces acceptable results, you should simply build the application on the OpenVMS I64 system (and on Alpha, if you wish) using the same qualifier.

For more information, see Section 4.8.4 and the "OpenVMS Floating-Point Arithmetic on the Intel® Itanium® Architecture" white paper. See the Related Documents section in the Preface for the web location of this white paper.

- Backup and restore functionality
- Operator interface look and feel
- System administration

## Assessing Needs for Porting Your Application

### 3.3 HP Porting Resources

### 3.3 HP Porting Resources

No single organization has expertise in all phases of a porting process. Therefore, HP has gathered a set of resources that can assist you through the process. These resources can provide everything from porting guides to teams of personnel with equipment that handles the entire porting task. For more information, consult your local HP representative and the Alpha Retain Trust services web site at:

[http://www.hp.com/products1/evolution/alpha\\_retaintrust/services.html](http://www.hp.com/products1/evolution/alpha_retaintrust/services.html)

A sampling of other resources available include the following:

- Developer and Solution Partner Program (DSPP)

<http://www.hp.com/go/dspp>

- Test Drive Program

<http://www.testdrive.hp.com/>

- HP Services

<http://www.hp.com/go/services/>

As the HP strategy on the Itanium architecture continues to develop, Sales Resources and Global Services will assess the impact to the long-term plans of HP customers and ISVs. HP will create and document template and custom service offerings to address generic and unique situations of each customer. HP will tailor each service offering to support your needs by making available the tools and resources needed to assist with crucial changes during the transition.

### 3.4 Additional Considerations

HP recognizes that not all layered software and middleware supplier products will port to OpenVMS I64 at first release. HP is committed to assist software vendors with information, hardware, support, and tools that will facilitate this process. Each vendor should monitor the progress of the supplier of their prerequisite software to ensure that the required components are in place for their development schedules.

HP intends to implement a tool suite on the existing OpenVMS Alpha platform that will make porting to the new OpenVMS I64 platform transparent for most developers in terms of the code base, commands, and process. Nothing will replace the testing that accompanies any quality product release. Rather, this tool suite will reduce the time required for product porting.

## Migrating Source Modules

This chapter describes the steps common to porting all applications. It then describes certain types of code and coding practices that may need to be changed before you can recompile your application.

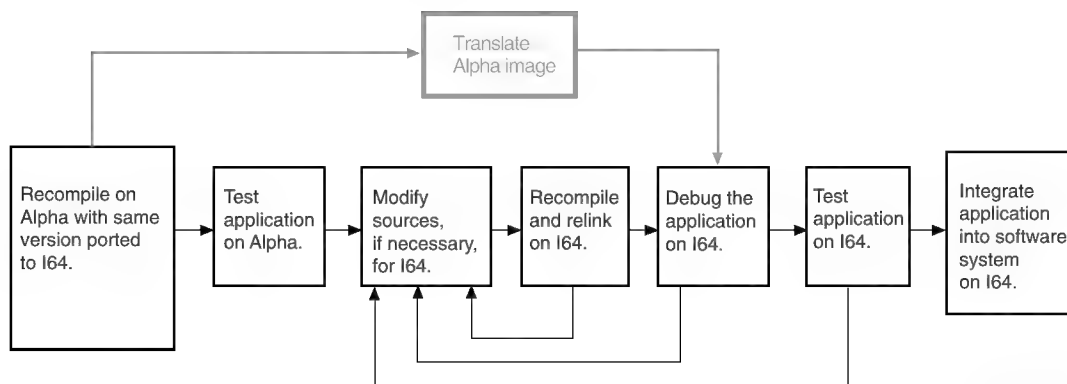
If you have thoroughly analyzed your code and planned the migration process, the final stage of migrating source modules should be fairly straightforward. You may be able to recompile or translate many programs with no change. Programs that do not recompile or translate directly will frequently need only straightforward changes to get them running on an I64 system. Certain types of code will require changes.

Migrating your application involves the following steps:

1. Setting up the migration environment.
2. Compiling your application on an Alpha system with the latest versions of the Alpha compilers.
3. Testing the application on an Alpha system to establish baselines for evaluating the migration.
4. Recompiling and relinking an application on an I64 system.
5. Debugging the migrated application.
6. Testing the migrated application.
7. Integrating the migrated application into a software system.

These steps, except for setting up the migration environment, are illustrated in Figure 4–1.

**Figure 4–1 Migrating an Application from Alpha to I64**



VM-1121A-AI

## Migrating Source Modules

If your source files possess any of the following characteristics, you may need to change them:

- Use VAX Macro-32 or Alpha Macro-64 assembly language
- Use conditions to produce the same application for different platforms from a common code base
- Rely on the Alpha architecture, such as the Alpha object code format
- Use floating point-data types
- Use threads
- Contain unaligned data
- Rely explicitly on the OpenVMS Alpha Calling Standard
- Use privileged interfaces or operate at inner access modes

For more information, see Section 4.8.

### 4.1 Setting Up the Migration Environment

The native I64 environment is comparable to the development environment on Alpha systems. Most of the program development tools that you are accustomed to using on OpenVMS Alpha are available. Some modifications have been made to these tools to account for the differences in the two architectures and the calling standard. Most of these modifications are invisible to users.

An important element of the I64 migration process is the availability of support from HP Services, which can provide help in modifying, debugging, and testing your application. For more information, see Section 3.3.

#### 4.1.1 Hardware

There are several issues to consider when planning what hardware you will need for your migration. To begin, consider what resources are required in your normal Alpha development environment:

- CPUs
- Disks
- Memory

In an I64 environment, the following configuration is recommended for this release:

- HP rx2600 Integrity Server with one or two processors
- 1 to 24 GB of memory
- One to four 36-GB disks
- DVD/CD-ROM drive

For more information, refer to the *HP OpenVMS Version 8.2 Release Notes*.



### 4.1.2 Software

To create an efficient migration environment, check the following elements:

- Migration tools

You need a compatible set of migration tools, including the following:

- Compilers
- Translation tools
  - HP OpenVMS Migration Software for Alpha to Integrity Servers and
  - TIE
- RTLs
- System libraries
- Include files for C programs

- Logical names

Logical names must be consistently defined to point to Alpha and I64 versions of tools and files.

- Compile and link procedures

These procedures may need to be adjusted for new tools and the new environment.

- Test tools

You need to port the Alpha test tools to OpenVMS I64, unless they are already ported. You also need test tools that are designed to test for OpenVMS I64 characteristics, such as changes to the OpenVMS Calling Standard that are specific to OpenVMS I64.

- Tools for maintaining sources and building images, such as:

- CMS (Code Management System)
- MMS (Module Management System)

#### **Native Alpha Development**

The standard development tools you have on Alpha are also available as native tools on I64 systems.

#### **Translation**

The software translator HP OpenVMS Migration Software for Alpha to Integrity Servers runs on both Alpha and I64 systems. The Translated Image Environment (TIE), which is required to run a translated image, is part of OpenVMS I64, so final testing of a translated image must either be done on an I64 system or at an I64 Migration Center. In general, the process of translating an Alpha image to run on an I64 system is straightforward, although you may have to modify your code somewhat to get it to translate without error.

## Migrating Source Modules

### 4.1 Setting Up the Migration Environment

#### 4.1.2.1 HP OpenVMS Migration Software for Alpha to Integrity Servers and Translated Image Environment (TIE)

The main tools for migrating Alpha user-mode images to OpenVMS I64 are a static translator and a run-time support environment:

- The HP OpenVMS Migration Software for Alpha to Integrity Servers is a utility that analyzes an Alpha image and creates a functionally equivalent translated image. Using HP OpenVMS Migration Software for Alpha to Integrity Servers, you will be able to do the following:
  - Determine whether an Alpha image is translatable.
  - Translate the Alpha image to an I64 image.
  - Identify specific incompatibilities with OpenVMS I64 within the image and, when appropriate, obtain information on how to correct those incompatibilities in the source files.
  - Identify ways to improve the run-time performance of the translated image.
- The Translated Image Environment (TIE) is an I64 shareable image that supports translated images at run time. TIE provides the translated image with an environment similar to OpenVMS Alpha and processes all interactions with the native I64 system. Items that TIE provides include:
  - Alpha instruction interpreter, which supports:
    - Execution of Alpha instructions (including instruction atomicity) that is similar to their execution on an Alpha system
    - Complex Alpha instructions, as subroutines
  - Alpha compatible exception handler
  - Jacket routines that allow communication between native and translated code
  - Emulated Alpha stack

TIE is invoked automatically for any translated image; you do not need to call it explicitly.

HP OpenVMS Migration Software for Alpha to Integrity Servers locates and translates as much Alpha code as possible into I64 code. TIE interprets any Alpha code that cannot be converted into I64 instructions; for example:

- Instructions that HP OpenVMS Migration Software for Alpha to Integrity Servers could not statically identify
- VAX (F\_, G\_ and D\_floating) floating-point operations

Since interpreting instructions is a slow process, HP OpenVMS Migration Software for Alpha to Integrity Servers attempts to find and translate as much Alpha code as possible to minimize the need for interpreting it at run time. A translated image runs at slower than a comparable native Alpha image, depending on how much Alpha code TIE needs to interpret.

Note that you cannot specify dynamic interpretation of a Alpha image on an I64 system. You must use HP OpenVMS Migration Software for Alpha to Integrity Servers to translate the image before it can run on OpenVMS I64.

Translating a Alpha image produces an image that runs as a native image on I64 hardware. The I64 image is not merely an interpreted or emulated version of the Alpha image, but contains I64 instructions that perform operations identical to those performed by the instructions in the original Alpha image. The I64 .EXE file also contains the original Alpha image in its entirety, which allows TIE to interpret any code that HP OpenVMS Migration Software for Alpha to Integrity Servers could not translate.

The HP OpenVMS Migration Software for Alpha to Integrity Servers's analysis capability also makes it useful for evaluating programs that you intend to recompile, rather than translate.

#### 4.1.3 Coexistence with Translated Images

Application components can be migrated from OpenVMS VAX and OpenVMS Alpha to OpenVMS I64 by binary translation. Because executable and shareable images can be translated on an individual basis, it is possible (even likely) to construct an environment in which native and translated images are mixed in a single image activation.

Translated images use the calling standard of their original implementation; that is, the argument lists of VAX translated images have the same binary format as they did on the VAX platform. Because the OpenVMS I64 calling standard differs from the OpenVMS VAX and OpenVMS Alpha calling standards, direct calls between native and translated images are not possible. Instead, such calls must be handled by interface routines (called jacket routines) that transform the arguments from one calling standard to the other.

The jacket routines are provided by TIE. In addition to argument transformation, TIE also handles the delivery of exceptions to translated code. TIE is loaded as an additional shareable image whenever a translated image (either main program or shareable image) is activated. The image activator interposes calls to jacket routines as necessary when it resolves references between native and translated images.

For native code to interoperate with translated code, it must meet the following requirements:

- All function descriptors must contain signature data. The signature data is used by the jacket routines to transform arguments and return values between their native and VAX or Alpha formats. Native I64 images contain function descriptors for both outbound calls and all entry points.
- All calls to function variables (that is, where the identity of the called function is contained in a variable) must be made through the library routine OTS\$CALL\_PROC. OTS\$CALL\_PROC determines at run time whether the function being called is native or translated and calls native code directly and translated code via the TIE.

These requirements are met by compiling code using the /TIE qualifier and linking it using the /NONATIVE\_ONLY qualifier. /TIE is supported by most OpenVMS I64 compilers including the Macro-32 compiler. (It is not supported on C++.) Both qualifiers must be given explicitly in the compile and LINK commands, since the defaults are /NOTIE and /NATIVE\_ONLY, respectively. Signatures are not supported by the I64 assembler; therefore, routines coded in I64 assembly language cannot be directly called from translated images. I64 assembly language can call translated code only with an explicit call to OTS\$CALL\_PROC.

## Migrating Source Modules

### 4.1 Setting Up the Migration Environment

Translated image support carries a minor performance penalty: all calls to function variables are made via `OTS$CALL_PROC`, at a cost of about 10 instructions for most cases. Performance-critical code should only be built `/TIE` and `/NONATIVE_ONLY` if interoperation with translated images is required.

An additional consideration applies to executable and shareable images that are linked `/NOSYSSHR`. Ordinarily, references to `OTS$CALL_PROC` are resolved by the shareable image `LIBOTS.EXE` and require no special attention. Linking an image `/NOSYSSHR` skips the search of the shareable image library and instead resolves external references from object modules in `STARLET.OLB`. `OTS$CALL_PROC` makes a reference to a data cell defined in the system symbol table, and if simply included from `STARLET.OLB` would result in an unresolved reference to `CTL$GQ_TIE_SYMVECT`.

Most images linked `/NOSYSSHR` are so built to avoid interaction with any external images, so calls to translated code are not an issue. For this reason, `STARLET.OLB` contains a subset version of `OTS$CALL_PROC` that does not support calls to translated images. This module, named `OTS$CALL_PROC_NATIVE`, is loaded by default, and images linked `/NOSYSSHR` by default cannot call out to translated code.

In addition, `STARLET.OLB` also contains the module `OTS$CALL_PROC`, which is the full version. It has no entry in the library symbol table and is only loaded by explicit reference. For the rare cases where images linked `/NOSYSSHR` need to be able to call out to translated shareable images, the full version of `OTS$CALL_PROC` must be explicitly included from `STARLET.OLB` and the image must also be linked against the system symbol table. Two options are required in the link command file:

- `SYS$LIBRARY:STARLET.OLB/INCLUDE=OTS$CALL_PROC` must be present in the input file list.
- The `/SYSEXE` qualifier must be included on the `LINK` command.

### 4.2 Compiling Applications on Alpha With Current Compiler Version

Before recompiling your code with a native I64 compiler, HP recommends that you first compile the code to run on Alpha with the latest version of the compiler. For example, if the application is written in Fortran, make sure that the application has been recompiled with Fortran Version 7.5 (Fortran 90) to run on Alpha. Fortran Version 7.5 is the version that was ported to I64.

Using the latest version of the compiler to recompile your application on OpenVMS Alpha might uncover problems that are attributable to the change in the compiler version only. For example, newer versions of compilers may enforce programming language standards that were previously ignored, exposing latent problems in your application code. Newer versions may also enforce changes to the standard that were not in effect when the earlier versions were created. Fixing such problems on OpenVMS Alpha simplifies porting that application to I64.

For a list of the native OpenVMS I64 compilers, see Section 5.1.

### 4.3 Testing Applications on Alpha for Baseline Information

The first step in testing is to establish baseline values for your application by running your test suite on the Alpha application. You can do this before or after you port your application to I64. You can then compare the results of these tests with the results of similar tests on an I64 system, as described in Section 4.6.

### 4.4 Recompiling and Relinking on an I64 System

In general, migrating your application involves repeated cycles of revising, compiling, linking, and debugging your code. During the process, you resolve all syntax and logic errors noted by the development tools. Syntax errors are usually simple to fix; logic errors typically require significant modifications to your code.

Your compile and link commands will likely require some changes, such as new compiler and linker switches. For more information about the compiler switches, see Chapter 6. If you are porting VAX MACRO code, refer to the *HP OpenVMS MACRO Compiler Porting and User's Guide*.

### 4.5 Debugging the Migrated Application

Once you have migrated your application to OpenVMS I64, you may have to debug it. This section describes the OpenVMS tools available for debugging your application.

For more information about these tools, see Chapter 5, Chapter 6, and the *HP OpenVMS Debugger Manual*.

#### 4.5.1 Debugging

The OpenVMS I64 operating system provides the following debuggers:

- OpenVMS Debugger

The OpenVMS Debugger is a symbolic debugger; that is, the debugger allows you to refer to program locations by the symbols you used for them in your program—the names of variables, routines, labels, and so on. You do not need to specify memory addresses or machine registers when referring to program locations.

The OpenVMS Debugger does not support debugging of translated images. However, you can debug a native application that uses a translated image.

- XDelta Debugger

The XDelta Debugger is an address location debugger; that is, the debugger requires you to refer to program locations by address location. This debugger is primarily used to debug programs that run in privileged processor mode or at an elevated interrupt level. The related Delta Debugger is not yet available.

The System-Code Debugger is a symbolic debugger that allows you to debug nonpageable code and device drivers running at any IPL.

#### 4.5.2 Analyzing System Crashes

OpenVMS provides two tools for analyzing system crashes: the System Dump Analyzer and the Crash Log Utility Extractor.

## Migrating Source Modules

### 4.5 Debugging the Migrated Application

#### 4.5.2.1 System Dump Analyzer

The System Dump Analyzer (SDA) utility on OpenVMS I64 systems is almost identical to the utility provided on OpenVMS Alpha systems. Many commands, qualifiers, and displays are identical, including several for accessing functions of the Crash Log Utility Extractor (CLUE) utility. Some displays have been adapted to show information specific to OpenVMS I64 systems.

To use SDA on an Alpha system, you must first familiarize yourself with the OpenVMS calling standard for Alpha systems. Similarly, to use SDA on an I64 system, you must familiarize yourself with the OpenVMS Calling Standard for I64 systems before you can decipher the pattern of a crash on the stack. For more information, refer to the *HP OpenVMS Calling Standard*.

#### 4.5.2.2 Crash Log Utility Extractor

The Crash Log Utility Extractor (CLUE) is a tool for recording a history of crash dumps and key parameters for each crash dump, and for extracting and summarizing key information. Unlike crash dumps, which are overwritten with each system failure and are available only for the most recent failure, the summary crash history file with a separate listing file for each failure is a permanent record of system failures.

### 4.6 Testing the Migrated Application

You must test your application to compare the functionality of the migrated version with that of the Alpha version.

The first step in testing is to establish baseline values for your application by running your test suite on the Alpha application, as described in Section 4.3.

Once your application is running on an I64 system, there are two types of tests you should apply:

- Standard tests used for the Alpha version of the application
- New tests to check specifically for problems due to the change in architecture

#### 4.6.1 Alpha Tests Ported to I64

Because the changes in your application are combined with use of a new architecture, testing your application after it is migrated to OpenVMS I64 is particularly important. Not only can the changes introduce errors into the application, but the new environment can bring out latent problems in the Alpha version.

Testing your migrated application involves the following steps:

1. Get a complete set of standard data for the application prior to the migration.
2. Migrate your Alpha test suite along with the application (if the tests are not already available on I64).
3. Validate the test suite on an I64 system.
4. Run the migrated tests on the migrated application.

Both regression tests and stress tests are useful here. Stress tests are important to test for platform differences in synchronization, particularly for applications that use multiple threads of execution.

### 4.6.2 New I64 Tests

Although your standard tests should go a long way toward verifying the function of the migrated application, you should add some tests that look at issues specific to the migration. Points to focus on include the following:

- Compiler differences
- Architectural differences
- Integration, such as modules written in different languages

### 4.6.3 Uncovering Latent Bugs

Despite your best efforts, you may encounter bugs that were in your program all along but that never caused a problem on an OpenVMS Alpha system.

For example, failure to initialize some variable in your program might have been benign on an Alpha system but could produce an arithmetic exception on an I64 system. The same could be true of moving between any other two architectures, because the available instructions and the way compilers optimize them is bound to change. There is no magic answer for bugs that have been in hiding, but you should test your programs after porting them and before making them available to other users.

## 4.7 Integrating the Migrated Application into a Software System

After you have migrated your application, check for problems that are caused by interactions with other software and that may have been introduced during the migration.

Sources of problems in interoperability can include the following:

- Alpha and I64 systems within an OpenVMS Cluster environment must use separate system disks. You must make sure that your application refers to the appropriate system disk.
- Image names

In a mixed architecture environment, be sure that your application refers to the correct version.

- Native Alpha and native I64 versions of an image have the same name.

If you attempt to run an Alpha image on an I64 system, the following error message is displayed:

```
$ run alpha_image.exe
%DCL-W-ACTIMAGE, error activating image alpha_image.exe
-CLI-E-IMGNAME, image file alpha_image.exe
-IMGACT-F-NOT_I64, image is not an HP OpenVMS Industry Standard 64 image
$
```

## 4.8 Modifying Certain Types of Code

Certain coding practices and types of code will require changes. The coding practices and types of code that may or will require changes are:

- Alpha Macro 64-Assembler code  
This code must be rewritten in another language.
- Code that has been conditionalized for running on Alpha or VAX systems  
This code must be revised to express an I64 condition.

## Migrating Source Modules

### 4.8 Modifying Certain Types of Code

- Code that uses OpenVMS system services that have dependencies on the Alpha architecture
- Code with other dependencies on the Alpha architecture
- Code that uses floating-point data types
- Code that uses a command definition file
- Code that uses threads, especially custom-written tasking or stack switching
- Privileged code

#### 4.8.1 Conditionalized Code

This section describes how to conditionalize OpenVMS code when migrating to I64. This code will be compiled for both Alpha and I64, or for VAX, Alpha, and I64. The symbol ALPHA, referred to in the following sections, is new. However, the symbol EVAX has not been eliminated. You do not need to replace EVAX with ALPHA but feel free to do so if convenient. The architecture symbols available for MACRO and BLISS are VAX, EVAX, ALPHA, and I64.

##### 4.8.1.1 MACRO Sources

For Macro-32 source files, the architecture symbols are in ARCH\_DEFS.MAR, which is a prefix file specified on the command line. On Alpha, ALPHA and EVAX equal 1 while VAX and I64 are undefined. On I64, I64 equals 1 while VAX, EVAX, and ALPHA are undefined.

The following example show how to conditionalize Macro-32 source code so that it can run on both Alpha and I64 systems.

For Alpha-specific code:

```
.IF DF ALPHA
    .
    .
    .
.ENDC
```

For I64-specific code:

```
.IF DF I64
    .
    .
    .
.ENDC
```

##### 4.8.1.2 BLISS Sources

For BLISS source files, either BLISS-32 or BLISS-64, the macros VAX, EVAX, ALPHA and I64 are defined in ARCH\_DEFS.REQ. On Alpha, EVAX and ALPHA equal 1 while VAX and I64 equal 0. On I64, I64 equals 1 while VAX, EVAX, and ALPHA equal 0. You must require ARCH\_DEFS.REQ to use these symbols in BLISS conditionals.

---

#### Note

---

The constructs %BLISS(*xxx*), %TARGET(*xxx*), and %HOST(*xxx*) are not recommended. Do not feel obligated to replace them, but feel free to do so if convenient.

---



Include the following statement in your source file:

```
REQUIRE 'SYS$LIBRARY:ARCH_DEFS';
```

Use the following statements in your source file to conditionalize code so that it can run on both Alpha and I64 systems.

For Alpha-specific code:

```
%if ALPHA %then
    .
    .
    .
%fi
```

For I64-specific code:

```
%if I64 %then
    .
    .
    .
%fi
```

#### 4.8.1.3 C Sources

For C source files, the symbols `__alpha`, `__ALPHA`, `__ia64`, and `__ia64__` are provided by the compilers on the appropriate platforms. Note that symbols could be defined on the compile command line but that is not the recommended method, nor is using `arch_defs.h`. Using `#ifdef` is considered the standard C programming practice.

For Alpha-specific code, use the following:

```
#ifdef __alpha
    .
    .
    .
#endif
```

For I64-specific code, use the following:

```
#ifdef __ia64
    .
    .
    .
#endif
```

#### 4.8.1.4 Existing Conditionalized Code

Existing conditionalized code must be examined to determine whether changes are required. Here is an example of BLISS code to consider.

```
%IF VAX %THEN
    VVV
    VVV
%FI

%IF EVAX %THEN
    aaa
    aaa
%FI
```

If the code is truly architecture specific and you are adding I64 code, then you would add the following case:

## Migrating Source Modules

### 4.8 Modifying Certain Types of Code

```
%IF I64 %THEN
    iii
    iii
%FI
```

However, if the existing VAX/EVAX conditionals really reflect 32 bits and not 64 bits or an “old” versus “new” OpenVMS convention (for example, a promoted data structure or different routine to call), then the following method for conditionalizing code might be more appropriate. That is because Alpha and I64 code are the same and 64-bit code need to be distinguished from the VAX code.

```
%IF VAX %THEN
    vvv
    vvv
%ELSE
    aaa
    aaa
%FI
```

#### 4.8.2 System Services With Alpha Architecture Dependencies

Certain system services that work well in applications on OpenVMS Alpha do not port successfully to I64. The following sections describe system services and their replacement services.

##### 4.8.2.1 SYS\$GOTO\_UNWIND

For OpenVMS Alpha, the SYS\$GOTO\_UNWIND system service accepts a 32-bit invocation context handle by reference. You must change instances of this system service to SYS\$GOTO\_UNWIND\_64, which accepts a 64-bit invocation context. Make sure to alter source code to allocate space for the 64-bit value. Also, different library routines return invocation context handles for OpenVMS I64. For more information, refer to the *HP OpenVMS Calling Standard*.

SYS\$GOTO\_UNWIND is most frequently used to support programming language features, so changes are mostly in compilers or run-time libraries. However, any direct use of SYS\$GOTO\_UNWIND requires change.

##### 4.8.2.2 SYS\$LKWSET and SYS\$LKWSET\_64

The SYS\$LKWSET and SYS\$LKWSET\_64 system services have been modified. For more information, see Section 4.8.9.

#### 4.8.3 Code With Other Dependencies on the Alpha Architecture

This section describes coding practices on Alpha that produce different results on I64 and may require changes to your application.

##### 4.8.3.1 Initialized Overlaid Program Sections

Initialized overlaid program sections are handled differently on I64 systems. On OpenVMS Alpha systems, different portions of an overlaid program section may be initialized by multiple modules. This is not allowed on OpenVMS I64 systems. For more information about this change in behavior, see Section 5.3.1.2.

#### 4.8.3.2 Condition Handlers Use of SS\$\_HPARITH

On OpenVMS Alpha, SS\$\_HPARITH is signaled for a number of arithmetic error conditions. On OpenVMS I64, SS\$\_HPARITH is never signaled for arithmetic error conditions; instead, the more specialized SS\$\_FLTINV and SS\$\_FLTDIV error codes are signaled on OpenVMS I64.

Update condition handlers to detect these more specialized error codes. In order to keep code common for both architectures, wherever the code refers to SS\$\_HPARITH, extend it for OpenVMS I64 to also consider SS\$\_FLTINV and SS\$\_FLTDIV.

#### 4.8.3.3 Mechanism Array Data Structure

The mechanism array data structure on OpenVMS I64 is very different from the one on OpenVMS Alpha. The return status code RETVAL has been extended to represent the return status register on both Alpha and I64 platforms. For more information, refer to the *HP OpenVMS Calling Standard*.

#### 4.8.3.4 Reliance on Alpha Object File Format

If your code relies on the layout of Alpha object files, you will need to modify it, because the object file format produced on OpenVMS I64 systems is different.

The object file format conforms to the 64-bit version of the executable and linkable format (ELF), as described in the *System V Application Binary Interface* draft of 24 April 2001. This document, published by Caldera, is available on their web site at:

<http://www.caldera.com/developers/gabi>

The object file format also conforms to the I64 specific extensions described in the *Intel® Itanium® Processor-specific Application Binary Interface (ABI)*, May 2001 edition (document number 245270-003). Extensions and restrictions, necessary to support object file and image file features that are specific to the OpenVMS operating system, will be published in a future release.

The portion of an image which is used by the debugger conforms to the DWARF Version 3 industry standard, which is available at the following location:

<http://www.eagercon.com/dwarf/dwarf3std.htm>

The debug symbol table representation on OpenVMS I64 is the industry-standard DWARF debug symbol table format described at this location. HP extensions to the DWARF Version 3 format will be published in a future release.

### 4.8.4 Code that Uses Floating-Point Data Types

OpenVMS Alpha supports VAX floating-point data types and IEEE floating point data types in hardware. OpenVMS I64 supports IEEE floating-point in hardware and VAX floating-point data types in software.

Most of the OpenVMS I64 compilers provide the /FLOAT=D\_FLOAT and /FLOAT=G\_FLOAT qualifiers to enable you to produce VAX floating-point data types. If you do not specify one of these qualifiers, IEEE floating-point data types will be used.

to specify a default floating-point data types using the I64 BASIC compiler, you use the /REAL\_SIZE qualifer.The possible values that can be specified are SINGLE (Ffloat), DOUBLE (Dfloat), GFLOAT, SFLOAT, TFLOAT, and XFLOAT.

## Migrating Source Modules

### 4.8 Modifying Certain Types of Code

You can test an application's behavior with IEEE floating-point values on Alpha by compiling it with an IEEE qualifier on OpenVMS Alpha. If that produces acceptable results, you can build the application on an I64 system using the same qualifier.

When you compile an OpenVMS application that specifies an option to use VAX floating-point on I64, the compiler automatically generates code for converting floating-point formats. Whenever the application performs a sequence of arithmetic operations, this code does the following:

1. Converts VAX floating-point formats to either IEEE single or IEEE double floating-point formats, as appropriate for the length.
2. Performs arithmetic operations in IEEE floating-point arithmetic.
3. Converts the resulting data from IEEE formats back to VAX formats.

Where no arithmetic operations are performed (VAX float fetches followed by stores), conversions do not occur. The code handles such situations as moves.

In a few cases, arithmetic calculations might have different results because of the following differences between VAX and IEEE formats:

- Values of numbers represented
- Rounding rules
- Exception behavior

These differences might cause problems for certain applications.

For more information about the differences between floating-point data types on OpenVMS Alpha and OpenVMS I64 and how these differences might affect ported applications, see Chapter 5 and refer to the “OpenVMS Floating-Point Arithmetic on the Intel® Itanium® Architecture” white paper. See the Related Documents section in the Preface for the web location of this white paper.

---

#### Note

---

Since the floating-point white paper was written, the default `/IEEE_MODE` has changed from `FAST` to `DENORM_RESULTS`. This means that, by default, floating-point operations might silently generate values that print as Infinity or Nan (the industry-standard behavior) instead of issuing a fatal run-time error as they would using VAX format float or `/IEEE_MODE=FAST`. Also, the smallest-magnitude nonzero value in this mode is much smaller because results are permitted to enter the denormal range instead of being flushed to zero as soon as the value is too small to represent with normalization. This default is the same default established by I64 at process startup.

---

#### 4.8.4.1 LIB\$WAIT Problem and Solution

The use of `LIB$WAIT` in code ported to OpenVMS I64 can cause unexpected results. An example in C follows.

```
float wait_time = 2.0;
lib$wait(&wait_time);
```

On OpenVMS I64 systems, this code sends an `S_FLOATING` into `LIB$WAIT`. `LIB$WAIT` expects an `F_FLOATING`, and gives a `FLTINV` exception.

LIB\$WAIT can accept three arguments. The previous code sequence can be rewritten with LIB\$WAIT using three arguments to run correctly on both I64 and Alpha systems. The following revised code works correctly when compiled without the /FLOAT qualifier:

```
#ifdef __ia64
    int float_type = 4; /* use S_FLOAT for I64 */
#else
    int float_type = 0; /* use F_FLOAT for Alpha */
#endif
float wait_time = 2.0;
lib$wait(&wait_time,0,&float_type);
```

A better coding method is to include LIBWAITDEF (from SYS\$STARLET\_C.TLB) in your application and then specify the floating point data types by name. This code can be maintained more easily.

LIBWAITDEF includes the following symbols:

- LIB\$K\_VAX\_F
- LIB\$K\_VAX\_D
- LIB\$K\_VAX\_G
- LIB\$K\_VAX\_H
- LIB\$K\_IEEE\_S
- LIB\$K\_IEEE\_T

The following example shows how to include libwaitdef.h in the code and how to specify the floating-point data type names. This example also assumes that the program is not compiled with /FLOAT.

```
#include <libwaitdef.h>
.
.
.
#ifdef __ia64
    int float_type = LIB$K_IEEE_S; /* use S_FLOAT for IPF */
#else
    int float_type = LIB$K_VAX_F; /* use F_FLOAT for Alpha */
#endif
float wait_time = 2.0;
lib$wait(&wait_time,0,&float_type);
```

#### 4.8.5 Incorrect Command Table Declaration

An incorrect declaration of a command table in code ported to OpenVMS I64 can cause unexpected results. For example, for an application, the Command Definition Utility is used to create an object module from a CLD file. The application then calls CLI\$DCL\_PARSE to parse command lines. CLI\$DCL\_PARSE may fail with:

```
%CLI-E-INTTAB, command tables have invalid format - see documentation
```

The code must be modified so that the command table is defined as an external data object.

For example, if in an application, on VAX and Alpha, the command table (DFSCP\_CLD) is incorrectly declared in a BLISS module as:

```
EXTERNAL ROUTINE DFSCP_CLD
```

## Migrating Source Modules

### 4.8 Modifying Certain Types of Code

This should be changed to

```
EXTERNAL DFSCP_CLD
```

Or if it was in a FORTRAN module incorrectly declared as:

```
EXTERNAL DFSCP_CLD
```

then it should be changed to

```
INTEGER DFSCP_CLD  
CDEC$ ATTRIBUTES EXTERN :: DFSCP_CLD
```

Similarly, in an application written in C, if the command tables previously were defined as follows:

```
int jams_master_cmd();
```

The code should be changed to be an external reference:

```
extern void* jams_master_cmd;
```

The changed, correct declaration works on all platforms: VAX, Alpha or I64.

#### 4.8.6 Code that Uses Threads

OpenVMS I64 supports all the thread interfaces that have been supported on OpenVMS since thread support was first introduced. Most OpenVMS Alpha code that uses threads can be ported to OpenVMS I64 without change. This section describes the exceptions. The major porting issue for code that uses threads is the usage of stack space. I64 code uses much more stack space than does equivalent Alpha code. Therefore, a threaded program that works on Alpha might get stack overflow failures on I64.

The default stack size is larger on OpenVMS I64 to help alleviate overflow problems. If the application requests a specific stack size, then the change to the default will be irrelevant, and the application source code might need to be changed to request more stack. HP recommends starting with an increase of three 8-Kb pages (24576 bytes).

Another side effect of the increased stack size requirement is increased demand on the P0 address region. Thread stacks are allocated from the P0 heap. Larger stacks might cause the process to exceed its memory quotas. In an extreme case, the P0 region could fill completely, in which case the process might need to reduce the number of threads in use concurrently (or make other changes to lessen the demand for P0 memory).

HP recommends that you familiarize yourself with the most recent improvements to thread support in OpenVMS, as documented in the *HP OpenVMS Version 8.2 Release Notes*. One change to the POSIX Threads C language header file `PTHREAD_EXCEPTION.H` caused problems when porting an application that relied on its former behavior.

The DCL command `THREADCP` is not supported on OpenVMS I64. For OpenVMS I64, the DCL commands `SET IMAGE` and `SHOW IMAGE` can be used to check and modify the state of threads-related image header flags, similar to the `THREADCP` command on OpenVMS Alpha. For details, refer to the *HP OpenVMS DCL Dictionary*.

The `THREADCP` command is documented in the *Guide to the POSIX Threads Library*.

If you want to change the setting of threads-related image flags, you need to use the new command SET IMAGE. For example:

```
$ SET IMAGE/FLAGS=(MKTHREADS,UPCALLS) FIRST.EXE
```

#### 4.8.6.1 Thread Routines `cma_delay` and `cma_time_get_expiration`

Two legacy threads API library routines, `cma_delay` and `cma_time_get_expiration`, accept a floating point format parameter using the VAX F FLOAT format. Any application modules that call either of these routines must be compiled with either the `/FLOAT=D_FLOAT` or the `/FLOAT=G_FLOAT` qualifier to get VAX F FLOAT support. (However, if your application also uses double precision binary data, then you must use the `/FLOAT=G_FLOAT` qualifier.) For more information about floating point support, consult your compiler's documentation.

If a C language module (which uses either `cma_delay` or `cma_time_get_expiration`) is compiled by mistake in an IEEE floating-point mode, a compiler warning similar to the following will be displayed:

```
cma_delay (  
%CC-W-LONGEXTERN, The external identifier name exceeds 31  
characters; truncated to "CMA_DELAY_NEEDS_VAX_FLOAT_____".
```

If an object file which triggered such a warning is linked, the linker will display an undefined-symbol message for this symbol. (If a linker-produced image was subsequently executed, the calls to these routines would fail with an ACCVIO.) These compiler and linker diagnostics are intended to alert you to the fact that these CMA routines require the use of VAX format floating-point values, and that the compilation was done in a manner that does not satisfy this requirement.

---

#### Note

These routines are no longer documented in user documentation but are still supported for use in legacy applications.

---

#### 4.8.7 Code With Unaligned Data

HP recommends that you align your data naturally to achieve optimal performance for data referencing. Unaligned data can seriously degrade performance on both OpenVMS Alpha and OpenVMS I64.

Data is **naturally aligned** when its address is an integral multiple of the size of the data in bytes. For example, a longword is naturally aligned at any address that is a multiple of 4, and a quadword is naturally aligned at any address that is a multiple of 8. A structure is naturally aligned when all its members are naturally aligned.

Because natural alignment is not always possible, OpenVMS I64 systems provide help to manage the impact of unaligned data references. Alpha and I64 compilers automatically correct most potential alignment problems and flag others.

In addition to performance degradation, unaligned shared data can cause a program to execute incorrectly. Therefore, you must align shared data naturally. Shared data might be between threads of a single process, between a process and ASTs, or between several processes in a global section.

## Migrating Source Modules

### 4.8 Modifying Certain Types of Code

#### Finding the Problem

To find instances of unaligned data, you can use a qualifier provided by most I64 compilers that allows the compiler to report compile-time references to unaligned data.

**Table 4–1 Compiler Switches for Reporting Compile-Time Reference**

Compiler	Switch
BLISS	/CHECK=ALIGNMENT
C	/WARN=ENABLE=ALIGNMENT
Fortran	/WARNING=ALIGNMENT
HP Pascal	/USAGE=PERFORMANCE

Additional assistance, such as an OpenVMS Debugger qualifier to reveal unaligned data at run time, is planned for a future release.

#### Eliminating the Problem

To eliminate unaligned data, you can use one or more of the following methods:

- Compile with natural alignment, or, when language semantics do not provide for this, move data to be naturally aligned. Where filler is inserted to ensure that data remains aligned, there is a penalty in increased memory size. A useful technique for ensuring naturally aligned data while conserving memory is to declare longer variables first.
- Use high-level-language instructions that force natural alignment within data structures.
- Align data items on quadword boundaries.

---

#### Note

Software that is converted to natural alignment may be incompatible with other software that is running translated in the same OpenVMS Cluster environment or over a network. For example:

- An existing file format may specify records with unaligned data.
- A translated image may pass unaligned data to, or expect it from, a native image.

In such cases, you must adapt all parts of the application to expect the same type of data, either aligned or unaligned.

---

#### 4.8.8 Code that Relies on the OpenVMS Alpha Calling Standard

If your application relies explicitly on characteristics of the OpenVMS Alpha calling standard, you likely have to change it. The OpenVMS I64 calling standard is based on the Intel calling standard with some OpenVMS modifications. Significant differences introduced in the OpenVMS I64 calling standard include the following:

- No frame pointer (FP)
- Multiple stacks
- Only four registers preserved across calls



- Register numbers you are familiar with have changed

For more information, see Chapter 2.

### **4.8.9 Privileged Code**

This section describes categories of privileged code that require examination, and may require modifications.

#### **4.8.9.1 Use of SYS\$LKWSET and SYS\$LKWSET\_64**

If your application uses SYS\$LKWSET or SYS\$LKWSET\_64 to lock itself into memory, and your application does not run on VAX systems, consider replacing these calls with calls to the new (as of OpenVMS Version 8.2) LIB\$LOCK\_IMAGE RTL routine. Similarly, replace the SYS\$ULWSET and SYS\$ULWSET\_64 calls with calls to the new LIB\$UNLOCK\_IMAGE RTL routine.

Programs that enter kernel mode and increase IPL to higher than 2 must lock program code and data in the working set. Locking code and data is necessary to avoid crashing the system with a PGFIPLHI bugcheck.

On VAX systems, typically only the code and data explicitly referenced by the program need to be locked. On Alpha, the code, data and linkage data referenced by the program need to be locked. On I64 systems, code, data, short data, and linker generated code need to be locked. To make porting easier and because the addresses of short data and linker generated data cannot be easily found within an image, changes have been made to the SYS\$LKWSET and SYS\$LKWSET\_64 system services on Alpha and I64.

As of OpenVMS Version 8.2, the SYS\$LKWSET and SYS\$LKWSET\_64 system services test the first address passed in. If this address is within an image, these services attempt to lock the entire image in the working set. If a successful status code is returned, the program can increase IPL to higher than 2 and without crashing the system with a PGFIPLHI bugcheck.

A counter is maintained within the internal OpenVMS image structures that counts the number of times the image has been successfully locked in the working set. The counter is incremented when locked and decremented when unlocked. When the counter becomes zero, the entire image is unlocked from the working set.

If your privileged program runs on Alpha and I64 and not VAX, you can remove all the code that finds the code, data and linkage data and locks these areas in the working set. You can replace this code with calls to LIB\$LOCK\_IMAGE and LIB\$UNLOCK\_IMAGE (available in OpenVMS Version 8.2). These routines are simpler to program correctly and make your code easier to understand and maintain.

If the program's image is too large to be locked in the working set, the status SS\$\_LKWSETFUL is returned. If you encounter this status, you can increase the user's working set quota. Otherwise, you can split the image into two parts, one that contains the user mode code and another sharable image that contains the kernel mode code. At the entry to a kernel mode routine, the routine should call LIB\$LOCK\_IMAGE to lock the entire image in the working set. Before exiting the kernel mode routine, the routine should call LIB\$UNLOCK\_IMAGE.

## Migrating Source Modules

### 4.8 Modifying Certain Types of Code

#### 4.8.9.2 Use of SYS\$LCKPAG and SYS\$LCKPAG\_64

If your application uses SYS\$LCKPAG or SYS\$LCKPAG\_64 to lock code in memory, examine your use of this service. On I64, this service will not lock the entire image in the working set.

It is likely that you intend to lock the image in the working set so your code can elevate IPL and execute without incurring a page fault. Refer to Section 4.8.9.1. See also the *HP OpenVMS RTL Library (LIB\$) Manual* for information about the LIB\$LOCK\_IMAGE and LIB\$UNLOCK\_IMAGE routines.

#### 4.8.9.3 Terminal Drivers

The interface for terminal class drivers on OpenVMS I64 is a call-based interface. This is a significant difference from the JSB-based interface on OpenVMS Alpha that uses registers to pass arguments.

The interface for OpenVMS I64 terminal class drivers is documented in the *OpenVMS Terminal Driver Port Class Interface for Itanium*. This document is available at the following location:

[http://www.hp.com/products1/evolution/alpha\\_retaintrust/openvms/resources](http://www.hp.com/products1/evolution/alpha_retaintrust/openvms/resources)

#### 4.8.9.4 Protected Image Sections

Protected image sections usually occur in shareable images which implement "User Written System Services."

These image sections are protected by software and hardware mechanisms to assure that an unprivileged application cannot compromise the integrity of these sections. Changes in hardware pages protection from VAX and Alpha to I64 have added some subtle restrictions which may require changes in the protected images.

As on VAX and Alpha, data sections that are writeable in privileged modes (kernel or exec) may be read by unprivileged (user) mode. The hardware protection for such pages does not allow execute access from any mode. Protected image sections which are linked as both writeable and executable are protected to allow inner mode read, write, and execute; no user mode access is allowed. As neither user mode access to inner mode writeable data, nor code being in writeable sections, is a common practice, it is felt that few applications are likely to require both in a single section.

While there were exceptions on VAX and Alpha, on I64, all writeable protected image sections on I64 are protected against user mode write. Protected images which intend to allow user write to protected image sections must use \$SETPRT/\$SETPRT\_64 to alter the page protection.

---

## OpenVMS I64 Development Environment

This chapter contains information about the OpenVMS I64 development environment, including:

- Native I64 compilers
- Other development tools
- Linker
- Debugger
- Librarian utility

### 5.1 Native I64 Compilers

For OpenVMS I64, native I64 compilers are available for the following languages:

- BASIC
- BLISS (available on the Freeware CD)
- C++
- COBOL V2.8
- Fortran V8.0 (Fortran 90)
- HP C V6.5
- HP Pascal
- Java
- VAX Macro-32

No support is provided to compile native Alpha Macro-64 assembler code to run on OpenVMS I64.

The I64 compilers provide command-line qualifiers for using VAX floating-point data types. For more information about these compilers, except for the VAX Macro-32 compiler, see Chapter 6.

#### 5.1.1 VAX MACRO-32 Compiler for OpenVMS I64

For new program development for OpenVMS I64, HP recommends the use of high-level languages. HP provides the VAX Macro-32 compiler for OpenVMS I64 to convert existing VAX MACRO code into machine code that runs on OpenVMS I64 systems.

Most VAX MACRO code can be compiled without any changes. The exceptions are:

- Programs that call routines that do not comply to the OpenVMS I64 calling standard.

## OpenVMS I64 Development Environment

### 5.1 Native I64 Compilers

- Programs that use the JSB instruction to call routines written in a language other than Macro-32.

For more information about porting VAX MACRO programs to OpenVMS I64 systems, refer to the *HP OpenVMS MACRO Compiler Porting and User's Guide*.

### 5.2 Other Development Tools

Several other tools in addition to the compilers are available to develop, debug, and deploy native I64 applications. These tools are summarized in Table 5–1.

**Table 5–1 OpenVMS Development Tools**

Tool	Description
OpenVMS Linker	The OpenVMS Linker accepts I64 object files to produce an I64 image. For more information about the OpenVMS Linker, see Section 5.3.
OpenVMS Debugger	The OpenVMS Debugger running on OpenVMS I64 has the same command interface as the current OpenVMS Alpha debugger. The graphical interface on OpenVMS Alpha systems will be available in a later release. For more information about the OpenVMS Debugger on OpenVMS I64, see Section 5.4.
XDelta Debugger	The XDelta Debugger is an address location debugger that is used for debugging programs that run in privileged processor mode or at an elevated interrupt priority level. The XDelta Debugger is available for this release but the related Delta Debugger is not yet available.
OpenVMS Librarian utility	The OpenVMS Librarian utility creates I64 libraries.
OpenVMS Message utility	The OpenVMS Message utility allows you to supplement the OpenVMS system messages with your own messages.
ANALYZE/IMAGE	The Analyze/Image utility can analyze I64 images.
ANALYZE/OBJECT	The Analyze/Object utility can analyze I64 objects.
DECset	DECset, a comprehensive set of development tools, includes the Language Sensitive Editor (LSE), the Digital Test Manager (DTM), Code Management System (CMS), and Module Management System (MMS). The two remaining DECset tools, Performance and Coverage Analyzer (PCA) and the Source Code Analyzer (SCA), will be available in a future release.
Command Definition utility	The Command Definition utility (CDU) enables application developers to create commands with a syntax similar to DCL commands.
System Dump Analyzer (SDA)	SDA has been extended to display information specific to OpenVMS I64 systems.
Crash Log Utility Extractor (CLUE)	CLUE is a tool for recording a history of crash dumps and key parameters for each crash dump, and for extracting and summarizing key information.

### 5.2.1 Translating Alpha Code

The HP OpenVMS Migration Software for Alpha to Integrity Servers is a binary translator for translating Alpha user code images to run on OpenVMS I64 systems. This tool is useful in cases where the source code is no longer available, or when the compiler is not available on I64, or when a third-party product has not been ported to OpenVMS I64. A third-party product's permission is required from the owner of the software.

The HP OpenVMS Migration Software for Alpha to Integrity Servers offers functionality similar to that of DECmigrate, a binary translator that was used to port OpenVMS VAX images to run on OpenVMS Alpha systems.

For more information, see Section 4.1.2.1.

## 5.3 Linking Modules

The purpose of the linker is to create images (files that contain binary code and data). The linker ported to OpenVMS I64 is different from the linker on OpenVMS VAX and Alpha systems because it accepts OpenVMS I64 object files and produces OpenVMS I64 images. As on OpenVMS VAX and Alpha systems, the primary type of image created is an **executable image**. This image can be activated at the DCL command line by issuing the RUN command.

The OpenVMS I64 Linker also creates **shareable images**. A shareable image is a collection of procedures and data that are exported via the symbol\_vector option, that can be called by executable images or other shareable images. Shareable images are included in an input file via a link operation that creates an executable or shareable image.

When linking modules, the primary type of input file to the OpenVMS I64 Linker is the **object file**. Object files are produced by language processors such as compilers or assemblers. Because the object files produced by these compilers are unique to the Intel Itanium architecture, some aspects of linking modules on OpenVMS I64 are different. In general, the OpenVMS I64 linker interface as well as functional capabilities (for example, symbol resolution, virtual memory allocation, and image initialization) are similar to those on OpenVMS VAX and Alpha systems. This section provides an overview of the differences as well as considerations you should review before linking programs on OpenVMS I64 systems. These include:

- Differences when linking on I64 systems as compared with linking on VAX and Alpha systems (Section 5.3.1)
- Expanded linker map file information for I64 (Section 5.3.2)
- New linker qualifiers and options (Section 5.3.3)
- Mixed-Case Arguments in Linker Options, Revisited (Section 5.3.4)

For details of these features and considerations, refer to the *HP OpenVMS Version 8.2 New Features and Documentation Overview*.

## OpenVMS I64 Development Environment

### 5.3 Linking Modules

#### 5.3.1 Differences When Linking on OpenVMS I64 Systems

Although linking on OpenVMS I64 systems is similar to linking on OpenVMS Alpha systems, some differences exist. The following qualifiers or options are ignored by the OpenVMS I64 linker:

- /REPLACE
- /SECTION\_BINDING
- /HEADER
- The per\_page keyword in /DEMAND\_ZERO=PER\_PAGE (the per\_page keyword will be supported, though with a slightly different meaning, in a future release)
- DZRO\_MIN
- ISD\_MAX

The following qualifiers and options are not allowed by the OpenVMS I64 Linker:

- /SYSTEM
- The file name keyword in /DEBUG=file\_spec
- The base address keyword must be null in CLUSTER=cluster\_name,base\_address . . .
- BASE=
- UNIVERSAL=

The following are new qualifiers supported by the OpenVMS I64 Linker:

- /BASE\_ADDRESS
- /SEGMENT\_ATTRIBUTE
- /FP\_MODE
- /EXPORT\_SYMBOL\_VECTOR
- /PUBLISH\_GLOBAL\_SYMBOLS
- The GROUP\_SECTIONS and SECTIONS\_DETAILS keywords for the /FULL command

Some of these qualifiers and options are described in the following sections. For more details on these qualifiers and options see the *HP OpenVMS Version 8.2 New Features and Documentation Overview*.

##### 5.3.1.1 No Based Clusters

Specifying a base address in a CLUSTER option is permitted on VAX and Alpha. On VAX even shareable images are allowed to contain based clusters, whereas on Alpha only main images are allowed to consist of such clusters. For I64 specifying a base address in a CLUSTER option is illegal for all images.

#### 5.3.1.2 Handling of Initialized Overlaid Program Sections on OpenVMS I64

On Alpha and VAX systems, initializations can be done to portions of an overlaid program section. Subsequent initializations to the same portions overwrite initializations from previous modules. The last initialization performed on any byte is used as the final one of that byte for the image being linked.

On I64 systems, the ELF (Executable and Linkable Format) object language does not implement the feature of the Alpha and VAX object language which allows the initialization of portions of sections. When an initialization is made, the entire section is initialized. Subsequent initializations of this section may be performed only if the non-zero portions match in value.

For example, the following condition produces different results on OpenVMS I64 systems than on Alpha systems:

Two program sections (simply termed 'sections' in ELF), each declaring two longwords, are overlaid. The first program section initializes the first longword and the second program section initializes the second longword with a non-zero value.

On Alpha systems, the linker is able to produce an image section with the first and second longword initialized. The VAX and Alpha object languages give the linker the section size and the Text Information Relocation (TIR) commands to initialize the section.

On I64 systems, the linker gets pre-initialized sections which contain initialization data from the compilers. The linker does not perform or know about the initializations since there are no TIR commands in ELF. The linker then produces a segment using the last processed section for initialization. That is, in the image either the first or second longword has a non-zero value, depending on the order in which the modules containing the sections were linked. On I64 systems, the linker reads the sections to be overlaid and checks for compatibility. Any two overlaid sections are compatible if they are identical in the non-zero values. If they are not compatible, the linker issues the following error:

```
%ILINK-E-INVOPRINI, incompatible multiple initializations for overlaid section
  section: <section name>
  module: <module name for first overlaid section>
  file: <file name for first overlaid section>
  module: <module name for second overlaid section>
  file: <file name for second overlaid section>
```

In the previous message, the linker lists the first module that contributes a non-zero initialization, and the first module with an incompatible initialization. Note that this is not a full list of all incompatible initializations; it is just the first one the linker encounters.

In the Program Section Synopsis of the linker map, each module with a non-zero initialization is flagged as "Initializing Contribution." Use this information to identify and resolve all the incompatible initializations. For further examples and more detail on the handling of initialized overlaid sections, see the *HP OpenVMS Version 8.2 New Features and Documentation Overview*.

#### 5.3.1.3 Behavior Difference When Linking ELF Common Symbols

The I64 linker behaves differently when ELF common symbols (also known as relaxed ref/def symbols) are linked selectively against an image that contains a definition for the same symbol. On Alpha, the linker incorrectly takes the definition from the relaxed ref/def symbol in your module. The I64 linker takes the definition from the shareable image.

## OpenVMS I64 Development Environment

### 5.3 Linking Modules

#### 5.3.2 Expanded Map File Information

Information in the linker map file has been expanded for OpenVMS I64 systems:

- On Alpha systems, the map section that was titled Object Module Synopsis has been renamed on I64 systems to Object and Image Synopsis, and contains expanded object and image information.
- On Alpha systems, the map section titled Image Section Synopsis has been divided into two sections for I64 systems: Cluster Synopsis and Image Segment Synopsis. The information has been moved to the appropriate section.
- In the Program Section Synopsis, the PIC and NOPIC attributes, invalid on I64 systems, have been removed.
- In the Symbol Cross Reference section, the designation for an external symbol has changed. The prefix or suffix used on Alpha was RX, meaning relocateable and external. However, the linker does not know whether an external symbol is relocatable or not. As a result, on I64 systems the prefix or suffix has been changed to X (external).
- In the Symbols By Value section, Keys for Special Characters have been changed or expanded. For example, UNIX weak symbols, (very likely only used by C++), designated by the UxWk key, are a new addition to OpenVMS I64.

For an example linker map illustrating this new information, refer to the *HP OpenVMS Version 8.2 New Features and Documentation Overview*.

#### 5.3.3 New Linker Qualifiers and Options for OpenVMS I64

Some new linker options and qualifiers have been added to support linking on OpenVMS I64 systems. This section describes these features.

##### 5.3.3.1 New /BASE\_ADDRESS Qualifier

A new qualifier, /BASE\_ADDRESS, is provided on I64 systems. The base address is the starting address that you want the linker to assign to an executable image. The purpose of this qualifier is to assign a virtual address for images which are not activated by the OpenVMS image activator, such as images used in the boot process. The OpenVMS image activator is free to ignore any linker assigned starting address. This qualifier is primarily used by system developers.

The /BASE\_ADDRESS qualifier does not replace the CLUSTER=[base-address] option, which is illegal on OpenVMS I64. See Section 5.3.1.1.

For more information on this qualifier, see the *HP OpenVMS Version 8.2 New Features and Documentation Overview*.

##### 5.3.3.2 New /SEGMENT\_ATTRIBUTE Qualifier

The OpenVMS I64 Linker provides a new /SEGMENT\_ATTRIBUTE which accepts two keywords: SHORT\_DATA=WRITE and DYNAMIC\_SEGMENT = P0 or P1. The DYNAMIC\_SEGMENT keyword is rarely needed. For more information see the *HP OpenVMS Version 8.2 New Features and Documentation Overview*.

The SHORT\_DATA=WRITE keyword allows you to combine read-only and read-write short data sections into a single segment, reclaiming up to 65,535 bytes of unused read-only space. When setting SHORT\_DATA to WRITE, your program may accidentally write to formerly read-only data. Therefore this qualifier is only recommended for users whose short data segment has reached the limit of 4 MB.



For more details on this qualifier, see the *HP OpenVMS Version 8.2 New Features and Documentation Overview*.

#### 5.3.3.3 New /FP\_MODE Qualifier

The OpenVMS I64 Linker determines the program's initial floating point mode using the floating point mode provided by the module that provides the main transfer address. Use the /FP\_MODE qualifier to set an initial floating point mode only if the module that provides the main transfer address does not provide an initial floating point mode. The /FP\_MODE qualifier will not override an initial floating point mode provided by the main transfer module. The OpenVMS I64 Linker accepts the following keywords to set the floating point mode:

- D\_FLOAT, G\_FLOAT—Sets VAX floating point modes.
- IEEE\_FLOAT[=ieee\_behavior]—Sets the IEEE floating point mode to the default or a specific behavior.

The OpenVMS I64 Linker accepts the following IEEE behavior keywords:

- FAST
- UNDERFLOW\_TO\_ZERO
- DENORM\_RESULTS (default)
- INEXACT

The OpenVMS I64 Linker also accepts a floating point mode behavior literal. For more information about the initial floating point mode, see the *HP OpenVMS Calling Standard*.

#### 5.3.3.4 New /EXPORT\_SYMBOL\_VECTOR and /PUBLISH\_GLOBAL\_SYMBOLS Qualifiers

The EXPORT\_SYMBOL\_VECTOR and PUBLISH\_GLOBAL\_SYMBOLS qualifiers were added to the linker to aid customers who are creating shareable images but did not know which symbols to export through the SYMBOL\_VECTOR option. This could be because they were porting an application from UNIX and were not familiar enough with the application to know which symbols to export, or because they were coding in C++, and were not able to know what the mangled names looked like.

The OpenVMS I64 Linker provides a new /PUBLISH\_GLOBAL\_SYMBOLS qualifier to mark an object module, so that all its global symbols can be exported in a symbol vector option. Additionally, the OpenVMS I64 Linker provides a new /EXPORT\_SYMBOL\_VECTOR qualifier to export a symbol vector option and to name the output file.

Both qualifiers are only accepted if the /SHAREABLE qualifier is present. /EXPORT\_SYMBOL\_VECTOR is a command line only qualifier. /PUBLISH\_GLOBAL\_SYMBOLS can be used in option files as well. The linker will warn, if there is an /EXPORT\_SYMBOL\_VECTOR qualifier but no /PUBLISH\_GLOBAL\_SYMBOLS qualifier was seen.

When /EXPORT\_SYMBOL\_VECTOR is present, only the option file is written, no image file is generated. The generated option file needs to be completed with GSMATCH information by the developer.

The /PUBLISH\_GLOBAL\_SYMBOLS qualifier is a positional qualifier and can be used with object files and libraries. It is compatible with the /INCLUDE and /SELECTIVE qualifiers, for example:

## OpenVMS I64 Development Environment

### 5.3 Linking Modules

```
$ link/SHARE public/PUBLISH,implementation/EXPORT=public
$ link/SHARE plib/LIBRARY/PUBLISH/INCLUDE=public/EXPORT=public
```

The `/EXPORT_SYMBOL_VECTOR` qualifier is a positional qualifier, which accepts an optional output file specification. If there is no file name specified the linker will use the name of the last input file and construct a file name, default file type is `".OPT"`.

For more detailed information on this qualifier, see the *HP OpenVMS Version 8.2 New Features and Documentation Overview*.

#### 5.3.3.5 New Alignments for the `PSECT_ATTRIBUTE` Option

The `PSECT_ATTRIBUTE` option now accepts integers 5, 6, 7, and 8 for the alignment attribute. The integers represent the byte alignment indicated as a power of 2. (For example, `2 ** 6` represents a 64-byte alignment.) The keyword `HEXA` (for hexadecimal word) was added for `2 ** 5` (that is, a 32-byte or 16-word alignment).

#### 5.3.3.6 New `GROUP_SECTIONS` and `SECTION_DETAILS` keywords for the `/FULL` Qualifier

The OpenVMS I64 Linker takes two keywords to the `/FULL` qualifier. The first keyword, `GROUP_SECTIONS`, prints all of the groups that were used in the map. Today the only compiler that takes advantage of groups is C++. Using this keyword with other languages has no effect.

When `/FULL=NOSECTION_DETAILS` is specified the OpenVMS I64 Linker suppresses zero length contributions in the Program Section Synopsis of the map. When the qualifier `/FULL` is used, it defaults to `/FULL=SECTION_DETAILS`, and a full linker map on VAX, Alpha, and I64 systems lists all the module contributions in the Program Section Synopsis.

### 5.3.4 Mixed-Case Arguments in Linker Options, Revisited

On OpenVMS I64 systems, names issued by compilers may be mixed case names. If you need to operate on mixed-case names in the options file (for example you have a library include statement and the module names are mixed-case) the linker already has an option to process the names in mixed-case, rather than using its default behavior (uppercasing all names). That option is the `CASE_SENSITIVE` option as shown below:

```
CASE_SENSITIVE=YES.
```

When the `CASE_SENSITIVE` option is set to `YES`, all characters to the right of the left-most equal sign (such as option arguments) have their case preserved. In other words, these characters are taken "as-is" without modification. This includes file names, module names, symbol names and keywords. To restore the linker's default behavior of upcasing the entire option line, specify the `CASE_SENSITIVE` option with the `NO` keyword as follows:

```
CASE_SENSITIVE=NO
```

Note that the `NO` keyword must appear in uppercase or it will not be recognized by the linker.

To maintain VAX and Alpha behavior it is recommended to switch to case sensitivity only when needed.

For more information and examples, refer to the *HP OpenVMS Version 8.2 New Features and Documentation Overview*.

## 5.4 Debugging Capabilities on OpenVMS I64 Systems

Several debuggers are provided on OpenVMS that allow for a wide range of debugging capabilities:

- OpenVMS Debugger (referred to simply as debugger)
- Delta (not available on OpenVMS I64)
- XDelta
- Source code debugger (SCD) (not available on OpenVMS I64)

The OpenVMS Debugger and XDelta debugger are available on OpenVMS 8.2. The following table shows the capabilities of these debuggers:

Category	OpenVMS Debugger	XDelta
Debugs operating systems	no	yes
Debugs applications	yes	no
Symbolic	yes	no
IPL greater than 0	no	yes
Process content	yes	no
User mode	yes	no†
Supervisor mode	no	no†
Exec mode	no	yes
Kernel mode	no	yes

†IPLs 0, 1, and 2 only

The following sections describe the capabilities of the OpenVMS Debugger and the XDelta debugger running on OpenVMS I64 systems. Further information about resolved problems, limitations, and known problems of the OpenVMS Debugger can be found in the *HP OpenVMS Version 8.2 Release Notes*.

### 5.4.1 OpenVMS Debugger

The debugger provided on OpenVMS I64 is different from the OpenVMS VAX and Alpha debugger. Although the OpenVMS I64 debugger shares a similar command interface and debugging functionality, some differences do exist. The following sections describe the capabilities of the OpenVMS I64 Debugger and assume a familiarity with the debugger on OpenVMS VAX and Alpha systems.

#### 5.4.1.1 Architecture Support

OpenVMS I64 debugger supports the following hardware registers:

- General registers R0 through R127
- Floating registers F0 through F127
- Branch registers B0 through B7
- A 64-bit predicate value named PRED, representing predicate registers P0 through P63
- Application registers: AR16 (RSC), AR17 (BSP), AR18 (BSPSTORE), AR19 (RNAT), AR25 (CSD), AR26 (SSD), AR32 (CCV), AR36 (UNAT), AR64 (PFS), AR65 (LC), AR66 (EC)

## OpenVMS I64 Development Environment

### 5.4 Debugging Capabilities on OpenVMS I64 Systems

- A program counter named PC, synthesized from the hardware IP register and the ri field of the PSR register
- Miscellaneous registers: CFM (current frame marker), UM (user mask), PSP (previous stack pointer), and IIPA (previously executed bundle address)

#### 5.4.1.2 Language Support

OpenVMS I64 Debugger supports programs written in the following languages:

- BLISS
- C
- C++ (limited)
- COBOL (limited)
- Macro-32
- Fortran
- Intel® Assembler (IAS)

Some issues exist when debugging supported languages on OpenVMS I64. For information about these issues and suggested workarounds, refer to the *HP OpenVMS Version 8.2 Release Notes*.

Support for C++ and COBOL in this release is limited. Generally, Debugger support for C++ is limited to programming constructs that are common to C, although support for basic C++ features such as references, classes, and class members is available.

Do not use the SHOW SYMBOL/ADDRESS and SHOW SYMBOL/TYPE commands for C++ data declarations. Specific examples of these and other problems are described in the *HP OpenVMS Version 8.2 Release Notes*.

Macro-32 is a compiled language on OpenVMS I64. Because the Itanium architecture has different hardware register usage than either Alpha or VAX, the IMACRO compiler must transform source code references to Alpha and VAX registers into a compatible register reference on Itanium. For complete information on register mapping used by Macro-32, refer to the *HP OpenVMS Version 8.2 Release Notes*.

#### 5.4.1.3 Functional Areas and Commands

The following functional areas and commands are available:

- Screen mode: source display, instruction display, output display
- Instruction decoding
- Breakpoints
- Watchpoints (non-static)
- Step (all forms)
- DEBUG/KEEP command (kept debugger configuration)
- RUN/DEBUG command (normal debugger configuration)
- Multithreaded programs
- CALL command
- Symbolic debugging of code in shareable images

## OpenVMS I64 Development Environment

### 5.4 Debugging Capabilities on OpenVMS I64 Systems

- Symbolization of all possible static data locations
- Values in floating-point registers are displayed as IEEE T-floating values
- Examines and deposits of floating-point variables
- EXAMINE/ASCII command

#### 5.4.1.4 Functionality Not Yet Ported

The following capabilities have not yet been ported to OpenVMS I64:

- DECwindows graphical user interface
- Heap analyzer
- Screen mode register view
- SHOW STACK command

#### 5.4.2 XDelta Debugger

In general, the XDelta Debugger for OpenVMS I64 systems behaves the same as XDelta on OpenVMS Alpha systems, with some restrictions. This section describes new capabilities added for OpenVMS and differences between XDelta on OpenVMS I64 and OpenVMS Alpha systems.

##### 5.4.2.1 XDelta Capabilities on OpenVMS I64

New capabilities have been added for OpenVMS I64 systems:

- XDelta supports the following Itanium registers:
  - General registers: R0 through R127
  - Floating registers: FP0 through FP127
  - Application registers: AR0 through AR127
  - Branch registers: BR0 through BR7
  - Control registers: CR0 through CR63
  - Miscellaneous registers: PC, PS, and CFM
  - Software implementation of Alpha registers
- ;D dump memory command
- ;T display interrupt stack command

For descriptions of these registers and capabilities in this release, refer to the *HP OpenVMS Version 8.2 New Features and Documentation Overview*. This document also describes the ;D and ;T commands, along with some restrictions.

##### 5.4.2.2 Differences Between XDelta on OpenVMS I64 and OpenVMS Alpha Systems

This section describes the differences in XDelta behavior on I64 and Alpha systems.

###### Interrupting a Running System

To interrupt a running system on I64, press Ctrl/P at the system console. Note that XDelta must have been previously loaded. When you press Ctrl/P, the system halts at the current PC and at the current IPL. There is no delay in waiting for the IPL to drop lower than 14 as on Alpha systems.

## OpenVMS I64 Development Environment

### 5.4 Debugging Capabilities on OpenVMS I64 Systems

#### Improved Symbolization

The X-registers are used by programmers to hold frequently used values, such as base addresses of images and modules. When displaying breakpoints and other address values, XDelta now prints these values relative to the nearest X-register value. Previously, only certain values were checked for proximity to X-register values.

## 5.5 OpenVMS I64 Librarian Utility

The Librarian utility on OpenVMS I64 provides the same features provided by Librarian on OpenVMS Alpha, with some changes and restrictions. This section describes what is unique about the I64 Librarian. For information about restrictions or other temporary conditions, refer to the *HP OpenVMS Version 8.2 Release Notes*.

### 5.5.1 Considerations When Using the I64 Librarian

You can use the DCL command LIBRARY (or Librarian LBR routines) to create libraries such as an I64 (ELF) object library, I64 (ELF) shareable image library, macro library, help library, and text library. You can maintain the modules in a library or display information about a library and its modules. The I64 Librarian cannot create or process Alpha and VAX objects and shareable images. Rather, the architecture for the Librarian is Intel Itanium.

There is no architecture switch for I64 libraries. The Librarian works with OpenVMS ELF object and image libraries when used with the following qualifiers:

- /OBJECT — Uses OpenVMS ELF object libraries (default).
- /SHARE — Uses OpenVMS ELF shareable image libraries.
- /CREATE — Creates an OpenVMS ELF library of an object or shareable image type, depending whether /OBJECT or /SHARE qualifier is specified.

The default library type created is an object library, if no OBJECT and SHARE qualifiers are specified.

### 5.5.2 Changes to the LBR\$ Routines

Two new library types for the LBR\$OPEN routine have been added on I64 systems:

LBR\$C\_TYP\_ELFOBJ (9) — Represents an ELF object library.

LBR\$C\_TYP\_ELFSHSTB (10) — Represents an ELF shareable image library.

In addition, the following library types for the LBR\$OPEN routine are not supported in the I64 Librarian. You cannot use the library types to create or open OpenVMS Alpha or VAX object and shareable image libraries:

LBR\$C\_TYP\_OBJ (1) — Represents a VAX object library.

LBR\$C\_TYP\_SHSTB (5) — Represents a VAX shareable image library.

LBR\$C\_TYP\_EOBJ (7) — Represents an Alpha object library.

LBR\$C\_TYP\_ESHSTB (8) — Represents an Alpha shareable image library.

### 5.5.3 I64 Library Format Handles UNIX-Style Weak Symbols

Due to the requirements of the Intel C++ compiler, the I64 library format has been expanded to accommodate new UNIX-style weak symbols. Multiple modules matching key names of new UNIX-style weak symbols can now exist in the same library. The Librarian ignores the OpenVMS-style weak symbol definitions as it has in the past.

UNIX-style weak symbol definitions behave in the same manner as weak transfer addresses on OpenVMS; that is, their definitions are tentative. If a definition of a stronger binding type is not seen during a link operation, the tentative definition is designated as the definitive definition.

#### 5.5.3.1 New ELF Type for Weak Symbols

A new Executable and Linkable Format (ELF) type was generated to distinguish between the two types of weak symbol definitions. For modules with ABI versions equal to 2 (the most common version used by compilers):

- Type STB\_WEAK represents the UNIX-style weak symbol (formerly, the OpenVMS-style weak symbol definition for ABI Version 1 ELF format)
- Type STB\_VMS\_WEAK represents the OpenVMS-style of weak symbol definition

The Librarian supports both the ELF ABI versions 1 and 2 of the object and image file formats within the same library.

---

#### Note

---

The new library format (Version 6.0) applies only to ELF object and shareable image libraries. Other libraries remain at the version 3.0 format. Applications that reference the library via the currently defined library services interface should not encounter any change in behavior.

---

#### 5.5.3.2 Version 6.0 Library Index Format

HP recommends using the new Version 6.0 libraries. Older (Version 3.0) libraries can be opened by the Library Services but cannot be not modified. They can be converted to Version 4.0 libraries, with some restrictions. For more information, refer to the *HP OpenVMS Version 8.2 New Features and Documentation Overview*.

#### 5.5.3.3 New Group-Section Symbols

Symbols may be relative to sections contained in an ELF entity called a **group**. These groups, and the symbols associated with them, behave in a similar fashion as the new UNIX-style weak symbol definitions; that is, they are tentative definitions. The librarian now allows multiple symbol definitions in the library's symbol name index.

#### 5.5.3.4 Current Library Limitation with Regard to Weak and Group Symbols

Library symbol entries are associated with the module within which they were defined. On OpenVMS I64, more than one module can define a UNIX-style weak or Group symbol. As such, the I64 Librarian must keep an ordered list of defining modules. This list allows a previous association to be restored should the higher precedence association be removed (you can find more information about precedence ordering rules in the *HP OpenVMS Version 8.2 Release Notes*.)

## OpenVMS I64 Development Environment

### 5.5 OpenVMS I64 Librarian Utility

Because of current limitations described in the *HP OpenVMS Version 8.2 Release Notes*, HP suggests that you only perform insert operations into the library only for modules that contain UNIX-style weak definitions. If you remove or replace modules in the library, you need to rebuild the library to make sure that the association of UNIX-style weak definitions is correct.



## Preparing to Port Applications

This chapter provides an overview of porting considerations related to the primary compilers that are available for OpenVMS I64. Table 6–1 maps the versions of Alpha compilers that were used to port to the I64 versions.

**Table 6–1 Mapping Alpha Compiler Versions to I64 Compiler Versions**

Compiler	On OpenVMS Alpha	On OpenVMS I64	For more information
BASIC	V1.6	V1.6	See Section 6.2.
BLISS	V1.11-004	V1.12-067	See Section 6.3.
COBOL	V2.8	V2.8	See Section 6.4.
Fortran 77	—	Not available <sup>1</sup>	See Section 6.5.2.
Fortran 90	V7.5	V8.0	See Section 6.5.
GNAT Pro Ada 95	Third party	—	See Section 6.1.
HP Ada 83	V3.5A	Not available <sup>1</sup>	—
HP C	V6.5	V7.1	See Section 6.6.
HP C++	V6.5	7.1	—
HP Pascal	V5.9	V5.9	—
Java	1.4.2	1.4.2-1	See Section 6.8.
Macro-32	V4.1-18	T1.0-77	See Section 6.9.
Macro-64	V1.2	Not available <sup>1</sup>	—

<sup>1</sup>Compiler is platform specific (for Alpha only).

Most of the OpenVMS I64 compilers share the following characteristics:

- They generate 32-bit code by default; you must use a compiler option to build a 64-bit program.
- The default is for IEEE floating-point data types instead of VAX floating-point data types.

The “OpenVMS Floating-Point Arithmetic on the Intel® Itanium® Architecture” white paper describes how I64 deals with floating-point issues. See the Related Documents section in the Preface for the web location of this white paper and other OpenVMS-to-Itanium architecture resources.

HP recommends that customers first compile their applications on OpenVMS Alpha with the versions of the compilers that are being ported to OpenVMS I64 in order to shake out any problems that might result from these newer compilers. (Some newer versions of compilers apply a stricter interpretation of existing compiler standards or enforce newer, stricter standards.) Once an application has

## Preparing to Port Applications

compiled, linked, and run without error on the new compiler on an OpenVMS Alpha system, it can be ported to OpenVMS I64.

### 6.1 Ada

Ada 95 is available on I64 in OpenVMS Version 8.2. Ada 83 is not supported on I64.

### 6.2 BASIC

The same BASIC is supported on both Alpha and I64. See the BASIC Release Notes for additional information on using BASIC on OpenVMS I64.

The BASIC compiler on OpenVMS Alpha defaults to /REAL\_SIZE=SINGLE (VAX F-float), and on OpenVMS I64 it defaults to /REAL\_SIZE=SFLOAT.

The BASIC compiler does not support the /IEEE\_MODE qualifier. The compiler and RTL set up the run-time environment on I64 to be similar in terms of exception handling and rounding to what is provided on Alpha.

### 6.3 BLISS Compiler

This section describes the differences between the OpenVMS Alpha and OpenVMS I64 BLISS compilers.

BLISS-32EN and BLISS-64EN are native compilers running on and generating code for OpenVMS for Alpha systems.

BLISS-32IN and BLISS-64IN are native compilers running on and generating code for OpenVMS I64 systems.

The BLISS-32 $xx$  compilers perform operations 32 bits wide (that is, BLISS values are longwords). The default width is 32 bits. In this chapter, they are collectively referred to as "the 32-bit compilers."

The BLISS-64 $xx$  compilers perform operations 64 bits wide (that is, BLISS values are quadwords). The default width is 64 bits. In this chapter, they are collectively referred to as "the 64-bit compilers".

The compilers are invoked using the following commands:

Platform	Compiler	Command
Alpha	BLISS-32EN	BLISS/A32 or BLISS
Alpha	BLISS-64EN	BLISS/A64
I64	BLISS-32IN	BLISS/I32 or BLISS
I64	BLISS-64IN	BLISS/I64

#### 6.3.1 BLISS File Types and File Location Defaults

This section discusses file types and output file location defaults for the BLISS compiler.

##### File Types

The default file type for object files for the OpenVMS compilers is .OBJ.

The default output file type for library files is .L32 for BLISS-32EN and BLISS-32IN, and .L64 for BLISS-64EN and BLISS-64IN. Library files are **not** compatible between dialects.

The search list for BLISS-32EN is as follows:

For source code: .B32E, .B32, .BLI  
For require files: .R32E, .R32, .REQ  
For library files: .L32E, .L32, .LIB

The search list for BLISS-64EN is as follows:

For source code: .B64E, .B64, .BLI  
For require files: .R64E, .R64, .REQ  
For library files: .L64E, .L64, .LIB

The search list for BLISS-32IN is as follows:

For source code: .B32I, .B32, .BLI  
For require files: .R32I, .R32, .REQ  
For library files: .L32I, .L32, .LIB

The search list for BLISS-64IN is as follows:

For source code: .B64I, .B64, .BLI  
For require files: .R64I, .R64, .REQ  
For library files: .L64I, .L64, .LIB

### Output File Location Defaults

For the OpenVMS compilers, regardless of whether they run on Alpha or I64, the location of the output files depends on where in the command line the output qualifier is found.

In both Alpha and I64 BLISS, if an output file qualifier, such as /OBJECT, /LIST, or /LIBRARY, is used after an input file specification and does not include an output file specification, the output file specification defaults to the device, directory, and file name of the immediately preceding input file. For example:

```
$ BLISS /A32 [FOO]BAR/OBJ      ! Puts BAR.OBJ in directory FOO
$ BLISS /I32 [FOO]BAR/OBJ      ! Puts BAR.OBJ in directory FOO
$
$ BLISS /A32 /OBJ [FOO]BAR      ! Puts BAR.OBJ in default directory
$ BLISS /I32 /OBJ [FOO]BAR      ! Puts BAR.OBJ in default directory
$
$ BLISS /A32 [FOO]BAR/OBJ=[ ]   ! Puts BAR.OBJ in default directory
$ BLISS /I32 [FOO]BAR/OBJ=[ ]   ! Puts BAR.OBJ in default directory
```

### 6.3.2 Alpha BLISS Features Not Available

This section describes Alpha BLISS features that are not supported by OpenVMS I64 BLISS.

#### Alpha BLISS Machine-Specific Builtins

Support for the following Alpha BLISS machine-specific builtins is no longer available:

CMP\_STORE\_LONG (replaced by CMP\_SWAP\_LONG)  
CMP\_STORE\_QUAD (replaced by CMP\_SWAP\_QUAD)  
CMPBGE  
DRAINT  
RPCC  
TRAPB  
WRITE\_MBX  
ZAP

## Preparing to Port Applications

### 6.3 BLISS Compiler

#### ZAPNOT

For information about `CMP_SWAP_LONG` and `CMP_SWAP_QUAD`, see `Compare` and `Swap` Builtin Functions.

#### Alpha BLISS PALcode Builtin Functions

Support for the following Alpha BLISS PALcode builtins is no longer available:

<code>CALL PAL</code>	<code>PAL MFPR PCBB</code>	<code>PAL MTPR SIRR</code>
<code>PAL BPT</code>	<code>PAL MFPR PRBR</code>	<code>PAL MTPR SSP</code>
<code>PAL BUGCHK</code>	<code>PAL MFPR PTBR</code>	<code>PAL MTPR TBIA</code>
<code>PAL CFLUSH</code>	<code>PAL MFPR SCBB</code>	<code>PAL MTPR TBIAP</code>
<code>PAL CHME</code>	<code>PAL MFPR SISR</code>	<code>PAL MTPR TBIS</code>
<code>PAL CHMK</code>	<code>PAL MFPR SSP</code>	<code>PAL MTPR TBISD</code>
<code>PAL CHMS</code>	<code>PAL MFPR TBCHK</code>	<code>PAL MTPR TBISI</code>
<code>PAL CHMU</code>	<code>PAL MFPR USP</code>	<code>PAL MTPR USP</code>
<code>PAL DRAINA</code>	<code>PAL MFPR VPTB</code>	<code>PAL MTPR VPTB</code>
<code>PAL HALT</code>	<code>PAL MFPR WHAMI</code>	<code>PAL PROBER</code>
<code>PAL GENTRAP</code>	<code>PAL MTPR ASTEN</code>	<code>PAL PROBEW</code>
<code>PAL IMB</code>	<code>PAL MTPR ASTSR</code>	<code>PAL RD PS</code>
<code>PAL LDQP</code>	<code>PAL MTPR DATFX</code>	<code>PAL READ UNQ</code>
<code>PAL MFPR ASN</code>	<code>PAL MTPR ESP</code>	<code>PAL RSCC</code>
<code>PAL MFPR ASTEN</code>	<code>PAL MTPR FEN</code>	<code>PAL STQP</code>
<code>PAL MFPR ASTSR</code>	<code>PAL MTPR IPIR</code>	<code>PAL SWPCTX</code>
<code>PAL MFPR ESP</code>	<code>PAL MTPR IPL</code>	<code>PAL SWASTEN</code>
<code>PAL MFPR FEN</code>	<code>PAL MTPR MCES</code>	<code>PAL WRITE UNQ</code>
<code>PAL MFPR IPL</code>	<code>PAL MTPR PRBR</code>	<code>PAL WR PS SW</code>
<code>PAL MFPR MCES</code>	<code>PAL MTPR SCBB</code>	<code>PAL MTPR PERFMON</code>

Macros are placed in `STARLET.REQ` for `PALCALL` builtins. OpenVMS supplies the supporting code. The privileged `CALL_PALs` call executive routines and the unprivileged `CALL_PALs` will call system services.

#### Alpha BLISS Register Names

The following registers (whose default I64 use is indicated in the following list) are not supported for naming in `REGISTER`, `GLOBAL REGISTER`, and `EXTERNAL REGISTER`, or as parameters to `LINKAGE` declarations.

<code>R0</code>	zero register
<code>R1</code>	global pointer
<code>R2</code>	volatile and GEM scratch register
<code>R12</code>	stack pointer
<code>R13</code>	thread pointer
<code>R14-R16</code>	volatile and GEM scratch registers
<code>R17-R18</code>	volatile scratch registers

#### INTERRUPT and EXCEPTION Linkages

`INTERRUPT` and `EXCEPTION` linkages are not supported.

#### BUILTIN Rn

You cannot specify an I64 register name to the `BUILTIN` keyword.

### 6.3.3 OpenVMS I64 BLISS Features

OpenVMS I64 BLISS provides only those existing Alpha BLISS features necessary to support OpenVMS I64. Although Alpha BLISS enabled support for features of operating systems other than OpenVMS, such functionality is not included in the I64 BLISS compiler.

#### OpenVMS I64 BLISS Builtins

Only those builtins necessary for the correct operation of OpenVMS I64 are supported by the BLISS I64 compilers.

**Common BLISS Builtins** The following existing common BLISS builtins are supported:

ABS	CH\$FIND_NOT_CH	CH\$WCHAR
ACTUALCOUNT	CH\$FIND_SUB	CH\$WCHAR_A
ACTUALPARAMETER	CH\$GEQ	MAX
ARGPTR	CH\$GTR	MAXA
BARRIER	CH\$LEQ	MAXU
CH\$ALLOCATION	CH\$LSS	MIN
CH\$A_RCHAR	CH\$MOVE	MINA
CH\$A_WCHAR	CH\$NEQ	MINU
CH\$COMPARE	CH\$PLUS	NULLPARAMETER
CH\$COPY	CH\$PTR	REF
CH\$DIFF	CH\$RCHAR	SETUNWIND
CH\$EQL	CH\$RCHAR_A	SIGN
CH\$FAIL	CH\$SIZE	SIGNAL
CH\$FILL	CH\$TRANSLATE	SIGNAL_STOP
CH\$FIND_CH	CH\$TRANSTABLE	

**RETURNADDRESS Builtin** A new builtin function RETURNADDRESS returns the PC of the caller's caller.

This builtin takes no arguments. The format is:

RETURNADDRESS()

#### Machine-Specific Builtins

The following Alpha BLISS machine-specific builtins are supported on the I64 BLISS compiler:

BARRIER  
ESTABLISH  
REVERT

ROT  
SLL  
SRA  
SRL  
UMULH

AdaWI

ADD_ATOMIC_LONG	AND_ATOMIC_LONG	OR_ATOMIC_LONG
ADD_ATOMIC_QUAD	AND_ATOMIC_QUAD	OR_ATOMIC_QUAD

The *xxx\_ATOMIC\_xxx* builtins no longer support the optional retry-count input argument. See ADD, AND, Builtin Functions for Atomic Operations for details.

TESTBITSSI	TESTBITCC	TESTBITCS
TESTBITCCI	TESTBITSS	TESTBITSC

TESTBIT<sub>xxx</sub> instructions no longer support the optional retry-count input argument or the optional success-flag output argument. See TESTBIT<sub>xxI</sub> and TESTBIT<sub>xx</sub> Builtin Functions for Atomic Operations for details.

ADDD	DIVD	MULD	SUBD	CMPD
ADDF	DIVF	MULF	SUBF	CMPF
ADDG	DIVG	MULG	SUBG	CMPG
ADDS	DIVS	MULS	SUBS	CMPS
ADDT	DIVT	MULT	SUBT	CMPT
CVTDF	CVTFD	CVTGD	CVTSF	CVTTD
CVTDG	CVTFG	CVTGF	CVTSI	CVTTG
CVTDI	CVTFI	CVTGI	CVTSL	CVTTI
CVTDL	CVTFL	CVTGL	CVTSQ	CVTTL
CVTDQ	CVTFQ	CVTGQ	CVTST	CVTTQ
CVTDT	CVTFS	CVTGT		CVTTS

## Preparing to Port Applications

### 6.3 BLISS Compiler

CVTID	CVTLD	CVTQD
CVTIF	CVTLF	CVTQF
CVTIG	CVTLG	CVTQG
CVTIS	CVTLS	CVTQS
CVTIT	CVTLT	CVTQT

CVTRDL	CVTRDQ
CVTRFL	CVTRFQ
CVTRGL	CVTRGQ
CVTRSL	CVTRSQ
CVTRTL	CVTRTQ

#### New Machine-Specific Builtins

A number of new builtins added to OpenVMS I64 BLISS provide access to single I64 instructions that can be used by the operating system.

**Builtins for Single I64 Instructions** Each name that is capitalized in the following list is a new builtin function that can be specified. The lowercase name in parenthesis is the actual I64 instruction executed. The arguments to these instructions (and therefore their associated BLISS builtin names) are detailed in the *Intel IA-64 Architecture Software Developer's Manual*.

BREAK	(break)	LOADRS	(loadrs)	RUM	(rum)
BREAK2	(break)*	PROBER	(probe.r)	SRLZD	(srlz.d)
FC	(fc)	PROBEW	(probe.w)	SRLZI	(srlz.i)
FLUSHRS	(flushrs)	PCTE	(ptc.e)	SSM	(ssm)
FWB	(fwb)	PCTG	(ptc.g)	SUM	(sum)
INVALAT	(invala)	PCTGA	(ptc.ga)	SYNCI	(sync.i)
ITCD	(itc.d)	PTCL	(ptc.l)	TAK	(tak)
ITCI	(itc.i)	PTRD	(ptr.d)	THASH	(thash)
ITRD	(itr.d)	PTRI	(ptr.i)	TPA	(tpa)
ITRI	(itr.i)	RSM	(rsm)	TTAG	(ttag)

---

#### Note

The BREAK2 builtin requires two parameters. The first parameter, which must be a compile-time literal, specifies the 21-bit immediate value of the BREAK instruction. The second parameter, can be any expression whose value is moved into IPF general register R17 just prior to executing the BREAK instruction.

---

**Access to Processor Registers** The OpenVMS I64 BLISS compiler provides builtin functions for access to read and write the many and varied processor registers in the IPF implementations. The builtin functions are as follows:

- GETREG
- SETREG
- GETREGIND
- SETREGIND

These builtins execute the mov.i instruction, which is detailed in the *Intel IA-64 Architecture Software Developer's Manual*. The two GET builtins return the value of the specified register.

To specify the register, a specially encoded integer constant is used, which is defined in an Intel C header file. See Appendix A of the *Intel IA-64 Architecture Software Developer's Manual* for the contents of this file.

**PALcode Builtins** Support for the following Alpha BLISS PALcode builtin functions has been retained:

PAL_INSQHIL	PAL_REMQHIL
PAL_INSQHILR	PAL_REMQHILR
PAL_INSQHIQ	PAL_REMQHIQ
PAL_INSQHIQR	PAL_REMQHIQR
PAL_INSQTIL	PAL_REMQTIL
PAL_INSQTILR	PAL_REMQTILR
PAL_INSQTIQ	PAL_REMQTIQ
PAL_INSQTIQR	PAL_REMQTIQR
PAL_INSQUEL	PAL_REMQUEL
PAL_INSQUEL_D	PAL_REMQUEL_D
PAL_INSQUEQ	PAL_REMQUEQ
PAL_INSQUEQ_D	PAL_REMQUEQ_D

Each of the 24 queue-manipulation PALcalls is implemented by BLISS as a call to an OpenVMS SYS\$PAL\_xxxx run-time routine.

### BLISSCALLG

The VAX idiom CALLG( .AP, ...) is replaced by an assembly routine BLISSCALLG(ARGPTR(), .RTN) for OpenVMS Alpha BLISS. This routine as defined for OpenVMS Alpha BLISS has been rewritten for the I64 architecture and is supported for OpenVMS I64 BLISS.

### I64 Registers

The I64 general registers, which can be named in REGISTER, GLOBAL REGISTER, and EXTERNAL REGISTER, and as parameters to LINKAGE declarations, are R3 through R11 and R19 through R31. In addition, parameter registers R32 through R39 can be named for parameters in LINKAGE declarations only.

Currently, there are no plans to support the naming of the I64 general registers R40 through R127.

Naming of any of the I64 Floating Point, Predicate, Branch and Application registers via the REGISTER, GLOBAL REGISTER, EXTERNAL REGISTER, and LINKAGE declarations is not provided.

A register conflict message is issued when a user request for a particular register cannot be satisfied.

### ALPHA\_REGISTER\_MAPPING Switch

A new module level switch, ALPHA\_REGISTER\_MAPPING, is provided for OpenVMS I64 BLISS.

This switch can be specified in either the MODULE declaration or a SWITCHES declaration. Use of this switch causes a remapping of Alpha register numbers to I64 register numbers as described in this section.

Any register number specified as part of a REGISTER, GLOBAL REGISTER, EXTERNAL REGISTER, or as parameters to GLOBAL, PRESERVE, NOPRESERVE or NOT USED in linkage declarations in the range of 0-31 are remapped according to the IMACRO mapping table as follows:

## Preparing to Port Applications

### 6.3 BLISS Compiler

0 = GEM TS REG K R8	16 = GEM TS REG K R14
1 = GEM TS REG K R9	17 = GEM TS REG K R15
2 = GEM TS REG K R28	18 = GEM TS REG K R16
3 = GEM TS REG K R3	19 = GEM TS REG K R17
4 = GEM TS REG K R4	20 = GEM TS REG K R18
5 = GEM TS REG K R5	21 = GEM TS REG K R19
6 = GEM TS REG K R6	22 = GEM TS REG K R22
7 = GEM TS REG K R7	23 = GEM TS REG K R23
8 = GEM TS REG K R26	24 = GEM TS REG K R24
9 = GEM TS REG K R27	25 = GEM TS REG K R25
10 = GEM TS REG K R10	26 = GEM TS REG K R0
11 = GEM TS REG K R11	27 = GEM TS REG K R0
12 = GEM TS REG K R30	28 = GEM TS REG K R0
13 = GEM TS REG K R31	29 = GEM TS REG K R29
14 = GEM TS REG K R20	30 = GEM TS REG K R12
15 = GEM TS REG K R21	31 = GEM TS REG K R0

The mappings for register numbers 16-20, 26-28, and 30-31 translate into registers that are considered invalid specifications for OpenVMS I64 BLISS (see Alpha BLISS Register Names and I64 Registers). Declarations that include any of these registers when ALPHA\_REGISTER\_MAPPING is specified will generate an error such as the following:

```
          r30 = 30,
          ^
.....
%BLS64-W-TEXT, Alpha register 30 cannot be declared, invalid mapping to
IPF register 12 at line number 9 in file ddd:[xxx]TESTALPHAREGMAP.BLI
```

Notice that the source line names register number 30 but the error text indicates register 12 is the problem. It is the translated register for 30, register 12, that is illegal to specify.

There is a special set of mappings for Alpha registers R16 through R21 if those registers are specified as linkage I/O parameters.

---

#### Note

---

For linkage I/O parameters *only*, the mappings for R16 through R21 are as follows:

16 = GEM TS REG K R32
17 = GEM TS REG K R33
18 = GEM TS REG K R34
19 = GEM TS REG K R35
20 = GEM TS REG K R36
21 = GEM TS REG K R37

---

#### ALPHA\_REGISTER\_MAPPING and "NOTUSED"

When ALPHA\_REGISTER\_MAPPING is specified, any Alpha register that maps to an IA64 scratch register and is specified as NOTUSED in a linkage declaration will be placed in the PRESERVE set.

This will cause the register to be saved on entry to the routine declaring it NOTUSED and restored on exit.



### **/ANNOTATIONS Qualifier**

The OpenVMS I64 BLISS compiler supports a new compilation qualifier, /ANNOTATIONS. This qualifier provides information in the source listing regarding optimizations that the compiler is making (or not making) during compilation.

The qualifier accepts the following keywords that reflect the different listing annotations:

- ALL
- NONE
- CODE — Used for annotations of machine code listing. Currently, only NOP instructions are annotated.
- DETAIL — Provides greater detail; used in conjunction with other keywords.

The remaining keywords reflect GEM optimizations:

INLINING  
LINKAGES  
LOOP\_TRANSFORMS  
LOOP\_UNROLLING  
PREFETCHING  
SHRINKWRAPPING  
SOFTWARE\_PIPELINING  
TAIL\_CALLS  
TAIL\_RECURSION

All keywords, with the exception of ALL and NONE, are negatable. The qualifier itself is also negatable. By default it is not present in the command line.

If the /ANNOTATIONS qualifier is specified without any parameters, the default is ALL.

### **/ALPHA\_REGISTER\_MAPPING Qualifier**

The OpenVMS I64 BLISS compiler supports a new compilation qualifier to enable ALPHA\_REGISTER\_MAPPING without having to modify the source. This is a positional qualifier. Specifying this qualifier on the compilation line for a module is equivalent to setting the ALPHA\_REGISTER\_MAPPING switch in the module header.

**/ALPHA\_REGISTER\_MAPPING Informationals** For OpenVMS I64 BLISS, three new informational messages have been added.

- If the /ALPHA\_REGISTER\_MAPPING qualifier was specified on the command line the following message is displayed:  

```
%BLS64-I-TEXT, Alpha Register Mapping enabled by the command line
```
- If the switch ALPHA\_REGISTER\_MAPPING is specified in the module header or as an argument to the SWITCH declaration the following message is displayed:

```
MODULE SIMPLE (MAIN=TEST, ALPHA_REGISTER_MAPPING)=  
.....^  
%BLS64-I-TEXT, Alpha Register Mapping enabled
```

## Preparing to Port Applications

### 6.3 BLISS Compiler

- If the switch NOALPHA\_REGISTER\_MAPPING is specified in the module header or as an argument to the SWITCH declaration the following message is displayed:

```
MODULE SIMPLE (MAIN=TEST, NOALPHA_REGISTER_MAPPING)=
.....^
%BLS64-I-TEXT, Alpha Register Mapping disabled
```

#### ADD, AND, Builtin Functions for Atomic Operations

The ADD\_ATOMIC\_XXXX, AND\_ATOMIC\_XXXX, and OR\_ATOMIC\_XXXX builtin functions for atomic updating of memory are supported by OpenVMS I64 BLISS. They are listed in Machine-Specific Builtins.

Because the I64 instructions to support these builtins waits until the operation succeeds, the optional retry-count input parameter has been eliminated. These builtins now have the form:

*option*\_ATOMIC\_*size*(*ptr*, *expr* [*old\_value*] ) !Optional output

where:

*option* can be AND, ADD, or OR.

*size* can be LONG or QUAD.

*ptr* must be a naturally aligned address.

*value* can be 0 (operation failed) or 1 (operation succeeded).

The operation is addition (or ANDing or ORing) of the expression EXPR to the data segment pointed to by PTR in an atomic fashion.

The optional output parameter OLD\_VALUE is set to the previous value of the data segment pointed to by PTR.

Any attempt to use the OpenVMS Alpha BLISS optional retry-count parameter results in a syntax error.

#### TESTBITxxI and TESTBITxx Builtin Functions for Atomic Operations

The TESTBITxxI and TESTBITxx builtin functions for atomic operations are supported by OpenVMS I64 BLISS. They are listed in Machine-Specific Builtins.

Because the I64 instruction to support these builtins waits until the operation succeeds, the optional input parameter retry-count and the optional output parameter success\_flag have been eliminated. These builtins now have the following form:

```
TESTBITxxx( field )
```

Any attempt to use the OpenVMS Alpha BLISS optional retry-count or success\_flag parameters results in a syntax error.

#### Granularity of Longword and Quadword Writes

OpenVMS I64 BLISS supports the /GRANULARITY=*keyword* qualifier, the switch DEFAULT\_GRANULARITY=*n*, and the data attribute GRANULARITY(*n*) as described in this section.

Users can control the granularity of stores and fetches by using the command line qualifier /GRANULARITY=*keyword*, the switch DEFAULT\_GRANULARITY=*n*, and the data attribute GRANULARITY(*n*).

The keyword in the command line qualifier must be either BYTE, LONGWORD, or QUADWORD. The parameter *n* must be either 0(byte), 2(longword) or 3(quadword).

When these are used together, the data attribute has the highest priority. The switch, when used in a SWITCHES declaration, sets the granularity of data declared after it within the same scope. The switch may also be used in the module header. The command line qualifier has the lowest priority.

#### Shift Builtin Functions

Builtin functions for shifts in a known direction are supported for OpenVMS I64 BLISS. They are listed in Machine-Specific Builtins. These functions are valid only for shift amounts in the range 0..%BPVAL-1.

#### Compare and Swap Builtin Functions

Both OpenVMS Alpha and OpenVMS I64 BLISS provide support for the following new compare and swap builtin functions:

- CMP\_SWAP\_LONG(addr, comparand, value)
- CMP\_SWAP\_QUAD(addr, comparand, value)

These functions do the following interlocked operations: compare the longword or quadword at addr with comparand and, if they are equal, store value at addr. They return an indicator of success (1) or failure (0).

#### I64-Specific Multimedia Instructions

There are no plans to support access to the I64-specific multimedia-type instructions.

#### Linkages

**CALL Linkage** The CALL linkage, as described in this section for OpenVMS Alpha BLISS, is also supported by OpenVMS I64 BLISS.

Routines compiled with a 32-bit compiler can call routines compiled with a 64-bit compiler and vice versa. Parameters are truncated when shortened, and sign-extended when lengthened.

By default, CALL linkages pass an argument count. This can be overridden using the NOCOUNT linkage option.

Although the arguments are passed in quadwords, the 32-bit compilers can “see” only the lower 32 bits.

**JSB Linkage** The OpenVMS I64 BLISS compilers have a JSB linkage type. Routines declared with the JSB linkage comply with the JSB rules developed for OpenVMS I64.

#### /[NO]TIE Qualifier

Support for this qualifier continues for OpenVMS I64. The default is /NOTIE.

TIE is used to enable the compiled code to be used in combination with translated images, either because the code might call into a translated image or might be called from a translated image.

In particular, TIE does the following:

- Causes the inclusion of procedure signature information in the compiled program. This may increase the size and possibly also the number of relocations processed during linking and image activation, but does not otherwise affect performance.

## Preparing to Port Applications

### 6.3 BLISS Compiler

- Causes calls to procedure values (sometimes called indirect or computed calls) to be compiled using a service routine (OTS\$CALL\_PROC); this routine determines whether the target procedure is native IPF code or in a translated image and proceeds accordingly.

#### **/ENVIRONMENT=([NO]FP) and ENVIRONMENT([NO]FP)**

The /ENVIRONMENT=([NO]FP) qualifier and the ENVIRONMENT([NO]FP) switch were provided for OpenVMS Alpha BLISS to cause the compiler to disable the use of floating point registers for certain integer division operations.

For OpenVMS I64 BLISS, the /ENVIRONMENT=NOFP command qualifier or ENVIRONMENT(NOFP) switch does not totally disable floating point because of the architectural features of I64. Instead, source code is still restricted not to have floating-point operations, but the generated code for certain operations (in particular, integer multiplication and division and the constructs that imply them) are restricted to using a small subset of the floating-point registers. Specifically, if this option is specified, the compiler is restricted to using f6-f11, and sets the ELF EF\_IA\_64\_REDUCEFP option described in the *Intel Itanium Processor-Specific Application Binary Interface*.

The /ENVIRONMENT=FP command qualifier and ENVIRONMENT(FP) switch are unaffected.

#### **Floating-Point Support**

The following sections discuss support for floating-point operations using the BLISS compiler.

**Floating-Point Builtin Functions** BLISS does not have a high level of support for floating-point numbers. The extent of the support involves the ability to create floating-point literals, and there are machine-specific builtins for floating-point arithmetic and conversion operations. For a complete list, see Machine-Specific Builtins.

None of the floating point builtin functions detect overflow, so they do not return a value.

**Floating-Point Literals** The floating-point literals supported by OpenVMS I64 BLISS are the same set supported by OpenVMS Alpha BLISS: %E, %D, %G, %S and %T.

**Floating-Point Registers** Direct use of the I64 floating-point registers is not supported.

**Calling Non-BLISS Routines with Floating-Point Parameters** It is possible to call standard non-BLISS routines that expect floating-point parameters passed by value and that return a floating-point or complex value.

The standard functions %FFLOAT, %DFLOAT, %GFLOAT, %SFLOAT and %TFLOAT are supported by OpenVMS I64 BLISS.

#### **New and Expanded Lexicals**

BLISS has added the following new compiler-state lexicals to support the OpenVMS I64 compilers: BLISS32I and BLISS64I.

- %BLISS now recognizes BLISS32V, BLISS32E, BLISS64E, BLISS32I and BLISS64I.
  - %BLISS(BLISS32) is true for all 32-bit BLISS compilers.
  - %BLISS(BLISS32V) is true only for VAX BLISS (BLISS-32).
  - %BLISS(BLISS32E) is true for all 32-bit Alpha compilers.

`%BLISS(BLISS64E)` is true for all 64-bit Alpha compilers.

`%BLISS(BLISS32I)` is true for all 32-bit I64 compilers.

`%BLISS(BLISS64I)` is true for all 64-bit I64 compilers.

- The lexicals `%BLISS32I` and `%BLISS64I` have been added. Their behavior parallels that of the new parameters to `%BLISS`.
- Support for the Intel Itanium architecture as a keyword to the `%HOST` and `%TARGET` lexicals has been added for OpenVMS I64 BLISS.

#### OpenVMS I64 BLISS Support for IPF Short Data Sections

The IPF calling standard requires that all global data objects with a size of 8 bytes or smaller be allocated in short data sections.

Short data sections can be addressed with an efficient code sequence that involves adding a 22-bit literal to the contents of the GP base register. This code sequence limits the combined size of all the short data sections. A linker error occurs if the total amount of data allocated to short data sections exceeds a size of  $2^{22}$  bytes. Compilers on IPF can use GP relative addressing when accessing short globals and short externals.

OpenVMS I64 BLISS exhibits new behavior for the PSECT attribute `GP_RELATIVE` and for the new PSECT attribute `SHORT`, which supports allocating short data sections.

Specifying the `GP_RELATIVE` keyword as a PSECT attribute causes PSECT to be labeled as containing short data so that the linker will allocate the PSECT close to the GP base address.

The syntax of the `SHORT` attribute is as follows:

```
"SHORT" "(" psect-name ")"
```

The following rules apply to the `SHORT` attribute:

- If the PSECT name in a `SHORT` attribute is not yet declared then its appearance in a `SHORT` attribute constitutes a declaration. The attributes of the PSECT containing the `SHORT` attribute become the attributes of the PSECT named in the `SHORT` attribute, except that the PSECT name declared in the `SHORT` attribute does not have the `SHORT` attribute and the PSECT name declared in the `SHORT` attribute does have the `GP_RELATIVE` attribute.
- If the PSECT name in a `SHORT` attribute has been previously declared then its attributes are not changed. A warning message is generated if the PSECT named in a `SHORT` attribute does not have the `GP_RELATIVE` attribute.
- If a data object with storage class `OWN`, `GLOBAL` or `PLIT` has a size of 8 or fewer bytes and the data object is specified to be allocated to a PSECT that includes the `SHORT` attribute, then that object is allocated to the PSECT named in the `SHORT` attribute. Note that this is a one-step process that is not recursive. If a short data object has its allocation PSECT renamed by the `SHORT` attribute, then the `SHORT` attribute of the renamed PSECT is not considered for any further renaming.
- Data objects with sizes larger than 8 bytes ignore the `SHORT` attribute.
- Data objects in the `CODE`, `INITIAL` and `LINKAGE` storage classes ignore the `SHORT` attribute, regardless of their size.

## Preparing to Port Applications

### 6.3 BLISS Compiler

- For the purposes of PSECT renaming by means of the SHORT attribute, the size of a PLIT object does not include the size of the count word that precedes the PLIT data.

**Example** The following example shows the use of PSECT in BLISS code.

```
PSECT

NODEFAULT = $GLOBAL SHORT$
  (READ,WRITE,NOEXECUTE,NOSHARE,NOPIC,CONCATENATE,LOCAL,ALIGN(3),
   GP_RELATIVE),

! The above declaration of $GLOBAL SHORT$ is not needed.  If the above
! declaration were deleted then the SHORT($GLOBAL_SHORT$) attribute in
! the following declaration would implicitly make an identical
! declaration of $GLOBAL_SHORT$.

GLOBAL = $GLOBAL$
  (READ,WRITE,NOEXECUTE,NOSHARE,NOPIC,CONCATENATE,LOCAL,ALIGN(3),
   SHORT($GLOBAL_SHORT$)),

NODEFAULT = MY_GLOBAL
  (READ,WRITE,NOEXECUTE,SHARE,NOPIC,CONCATENATE,LOCAL,ALIGN(3)),

PLIT = $PLIT$
  (READ,NOWRITE,NOEXECUTE,SHARE,NOPIC,CONCATENATE,GLOBAL,ALIGN(3),
   SHORT($PLIT_SHORT$));

GLOBAL
  X1,                      ! allocated in $GLOBAL_SHORT$
  Y1 : VECTOR[2, LONG],    ! allocated in $GLOBAL_SHORT$
  Z1 : VECTOR[3, LONG],    ! allocated in $GLOBAL$
  A1 : PSECT(MY_GLOBAL),   ! allocated in MY_GLOBAL
  B1 : VECTOR[3, LONG] PSECT(MY_GLOBAL), ! allocated in MY_GLOBAL
  C1 : VECTOR[3, LONG]
      PSECT($GLOBAL_SHORT$); ! allocated in $GLOBAL_SHORT$

PSECT GLOBAL = MY_GLOBAL;
! use MY_GLOBAL as default for both noshort/short

GLOBAL
  X2,                      ! allocated in MY_GLOBAL
  Y2 : VECTOR[2, LONG],    ! allocated in MY_GLOBAL
  Z2 : VECTOR[3, LONG],    ! allocated in MY_GLOBAL
  A2 : PSECT($GLOBAL$),    ! allocated in $GLOBAL_SHORT$
  B2 : VECTOR[3, LONG] PSECT($GLOBAL$); ! allocated in $GLOBAL$;

! Note that the allocations of A1, X2 and Y2 violate the calling
! standard rules.  These variables cannot be shared with other
! languages, such as C or C++.

PSECT GLOBAL = $GLOBAL$;
! back to using $GLOBAL/$GLOBAL_SHORT$ as default noshort/short

GLOBAL BIND
  P1 = UPLIT("abcdefghi"), ! allocated in $PLIT$
  P2 = PLIT("abcdefgh"),   ! allocated in $PLIT_SHORT$
  P3 = PSECT(GLOBAL) PLIT("AB"), ! allocated in $GLOBAL_SHORT$
  P4 = PSECT($PLIT_SHORT$)
      PLIT("abcdefghijklmn"), ! allocated in $PLIT_SHORT$
  P5 = PSECT(MY_GLOBAL) PLIT("AB"); ! allocated in MY_GLOBAL
```

---

#### Notes

- The allocations of A1, X2, Y2, and P5 violate the calling standard rules. These variables cannot be shared with other languages, such

as C or C++. They can be shared with modules written in BLISS and MACRO.

- The current OpenVMS I64 BLISS design does not support GP\_RELATIVE addressing mode on EXTERNAL variable references. However, the usual GENERAL addressing mode used by EXTERNAL variables correctly references a GP\_RELATIVE section. Currently, there are no plans to add an ADDRESSING\_MODE(GP\_RELATIVE) attribute to BLISS.
- 

## 6.4 COBOL

COBOL Version 2.8 is supported on both Alpha and I64. See the *COBOL Release Notes* for restrictions and known problems related to using COBOL on OpenVMS I64.

### 6.4.1 Floating-Point Arithmetic

The COBOL compiler on OpenVMS Alpha defaults to /FLOAT=D\_FLOAT. For I64, the default is /FLOAT=IEEE\_FLOAT.

The COBOL compiler does not support /IEEE\_MODE. The COBOL RTL sets up the run-time environment on I64 to be similar in terms of exception handling and rounding to what is provided in the COBOL run-time environment on Alpha.

The COBOL Release Notes and the white paper entitled “OpenVMS Floating-Point Arithmetic on the Intel® Itanium® Architecture” together describe how COBOL deals with floating-point issues on I64.

See the Related Documents section in the Preface for the web location of this white paper and other OpenVMS-to-Itanium architecture resources.

### 6.4.2 /ARCH and /OPTIMIZE=TUNE Qualifiers

For the sake of “compile-and-go” compatibility, Alpha values for the /ARCH and /OPTIMIZE=TUNE qualifiers are accepted by the COBOL compiler on I64. An informational message is displayed indicating that they are ignored.

I64 values for /ARCH and /OPTIMIZE=TUNE are defined in the COBOL compiler for development purposes only. Their behavior is unpredictable and they should not be used.

## 6.5 Fortran

Fortran 90 V8.0 is supported on OpenVMS I64. Fortran 77 is not supported on OpenVMS I64. (See Section 6.5.2 for details.)

HP Fortran for OpenVMS I64 systems provides the same command-line options and language features as HP Fortran for OpenVMS Alpha systems with a few exceptions. These exceptions are discussed in the following sections.

### 6.5.1 Floating-Point Arithmetic

The Fortran release notes and the white paper entitled “OpenVMS Floating-Point Arithmetic on the Intel® Itanium® Architecture” together describe how HP Fortran for I64 deals with floating-point issues.

See the Related Documents section in the Preface for the web location of this white paper and other OpenVMS-to-Itanium architecture resources.

## Preparing to Port Applications

### 6.5 Fortran

Highlights are summarized as follows:

- IEEE is the default floating-point datatype (that is, the default is `/FLOAT=IEEE_FLOAT`).

The OpenVMS Alpha compiler defaults to `/FLOAT=G_FLOAT`. On OpenVMS I64 systems, there is no hardware support for floating-point representations other than IEEE. The VAX floating-point formats are supported in the compiler by generating run-time code to convert to IEEE format before performing arithmetic operations, and then converting the IEEE result back to the appropriate VAX format. This imposes additional run-time overhead and some loss of accuracy compared to performing the operations in hardware on Alpha (and VAX).

The software support for the VAX formats is an important functional compatibility requirement for certain applications that need to deal with on-disk binary floating-point data, but its use should not be encouraged. This change is similar to the change in default from `/FLOAT=D_FLOAT` on VAX to `/FLOAT=G_FLOAT` on Alpha.

If at all possible, users should use `/FLOAT=IEEE_FLOAT` (the default) for the highest performance and accuracy.

Note that the changed `/FLOAT` default has implications for the use of `/CONVERT=NATIVE` (the default). This switch causes unformatted data to remain unconverted on input, on the assumption that it matches the prevailing floating-point datatype.

Files written from Fortran applications built with `/FLOAT=G_FLOAT/CONVERT=NATIVE` on Alpha (either explicitly `G_FLOAT` and `NATIVE` or defaulting either or both) will be in `G_FLOAT`. They can be read by I64 applications built with `/FLOAT=G_FLOAT/CONVERT=NATIVE` or `/FLOAT=IEEE/CONVERT=VAXG`. But they will not be readable by applications built with `/CONVERT=NATIVE`, as that will default to the `/FLOAT` type, which will have defaulted to IEEE.

- The `/IEEE_MODE` qualifier defaults to `/IEEE_MODE=DENORM_RESULTS`. This differs from Alpha, where the default is `/IEEE_MODE=FAST`. Despite the name, `/IEEE_MODE=FAST` does not have a significant effect on run-time performance on I64 (or on Alpha processors from EV6 onward).
- Users must choose one `/FLOAT` value and one `/IEEE_MODE` value and stick with it for the whole of their application. This is because mixed-mode applications do not (in general) work on OpenVMS I64 systems as a result of architectural differences in the hardware. This is a change from OpenVMS on Alpha systems, where mixed-mode applications work. In particular, per-routine/per-file/per-library settings of a mode do not work.

As on Alpha, `/IEEE_MODE` can only be set if the user has chosen or defaulted to `/FLOAT=IEEE`. The `IEEE_MODE` used for `G_FLOAT` and `D_FLOAT` is one the compiler has picked as appropriate for supporting the IEEE-format calculations which implement VAX-format support.

- If your code changes the floating-point exception mode, it is your responsibility to change it back on routine exit, including exits by means of exception handlers. Failure to do so might lead to unexpected results, since the compilers assume that no called routine changes the current mode. This is also a requirement on user-written libraries.



- Exception handlers are entered with the floating-point mode in effect at the time the exception was raised, not the mode with which the handler was compiled.

As specified in the white paper, *HP OpenVMS Floating-Point Arithmetic on the Itanium® Architecture*, the number, type and location of floating-point exceptions raised during the execution of an application on I64 may not be the same as on Alpha. This is particularly true for VAX-format floating-point. In that case exceptions are (in general) only raised on the convert back to VAX-format after the computation is over.

Three consequences may be of interest for users of VAX-format floating-point:

- Intermediate values (such as the "X/Y" in "X\*Y + X/Y" which could have raised an exception on Alpha will probably not raise an exception on I64.
- Values not converted back to VAX-format floating-point (such as the "X/Y" in "IF (X/Y > 0.0)") will probably not raise an exception on I64.
- OpenVMS Alpha and VAX Fortran applications do not report underflows for VAX-format floating-point operations unless you specifically enable underflow traps by compiling with the /CHECK=UNDERFLOW qualifier. The same is true on OpenVMS I64 systems, but with an important caveat.

Since all I64 floating-point operations are implemented using IEEE-format operations, enabling underflow traps with /CHECK=UNDERFLOW causes exceptions to be raised when values underflow the IEEE-format representation, not the VAX-format one.

This can result in an increased number of underflow exceptions seen with /CHECK=UNDERFLOW when compared with equivalent Alpha or VAX programs, as the computed values may be in the valid VAX-format range, but in the denormalized IEEE-format range.

### 6.5.2 Only the F90 Compiler is Supported

The F77 compiler, previously invoked with the /OLD\_F77 qualifier, is not available. Development is currently underway to provide the following functionality contained in the Alpha F77 compiler that is not available in the I64 and Alpha F90 compilers:

- CDD and FDML support, which were previously available only in the Alpha F77 compiler, have now been implemented in the I64 Fortran 90 compiler.
- Code previously compiled with the F77 compiler that does not compile with the F90 compiler should be reported through the problem reporting mechanism. Although it may not be possible to eliminate all differences, the HP development team has attempted to minimize these differences by enhancing the F90 compiler.

---

#### Note

The lack of support for /OLD\_F77 should not be confused with the /F77 qualifier, which indicates that the compiler uses FORTRAN-77 interpretation rules for those statements that have a meaning incompatible with FORTRAN-66, and which will be supported.

---

## Preparing to Port Applications

### 6.5 Fortran

#### 6.5.3 /ARCH and /OPTIMIZE=TUNE Qualifiers

For the sake of “compile-and-go” compatibility, Alpha values for the /ARCH and /OPTIMIZE=TUNE qualifiers are accepted on the compiler invocation command. An informational message is printed saying they are ignored.

I64 values for /ARCH and /OPTIMIZE=TUNE are defined in the OpenVMS I64 compiler for development purposes only. Their behavior is unpredictable and they should not be used.

## 6.6 HP C/ANSI C

HP C Version 7.1 is supported on I64 systems. Refer to the HP C documentation, including the HP C Release Notes for details and additional porting considerations.

#### 6.6.1 I64 Floating-Point Default

The native Alpha compiler defaults to /FLOAT=G\_FLOAT. For I64, the default is /FLOAT=IEEE\_FLOAT/IEEE=DENORM.

On OpenVMS I64, there is no hardware support for floating-point representations other than IEEE. The VAX floating point formats are supported in the compiler by generating run-time code to convert to IEEE format in order to perform arithmetic operations, and then convert the IEEE result back to the appropriate VAX format. This imposes additional run-time overhead and a possible loss of accuracy compared to performing the operations in hardware on the Alpha (and VAX). The software support for the VAX formats is an important functional compatibility requirement for certain applications that deal with on-disk binary floating-point data, but its use should not be encouraged by letting it remain the default. This change is similar to the change in default from /FLOAT=D\_FLOAT on VAX to /FLOAT=G\_FLOAT on Alpha.

Note also that the default /IEEE\_MODE has changed from FAST (on OpenVMS Alpha) to DENORM\_RESULTS on OpenVMS I64. This means that, by default, floating-point operations silently generate values that print as Infinity or Nan (the industry-standard behavior) instead of issuing a fatal run-time error as they would using VAX format float or /IEEE\_MODE=FAST. Also, the smallest-magnitude nonzero value in this mode is much smaller because results are permitted to enter the denormal range instead of being flushed to zero as soon as the value is too small to represent with normalization.

#### 6.6.2 Semantics of /IEEE\_MODE Qualifier

On OpenVMS Alpha, the /IEEE\_MODE qualifier generally has its greatest effect on the generated code of a compilation. When calls are made between functions compiled with different /IEEE\_MODE qualifiers, each function produces the /IEEE\_MODE behavior with which it was compiled.

On OpenVMS I64, the /IEEE\_MODE qualifier primarily affects only the setting of a hardware register at program startup. In general, the /IEEE\_MODE behavior for a given function is controlled by the /IEEE\_MODE option that was specified on the compilation that produced the main program. The compiler marks the object module of each compilation with the floating-point control options specified by the compile-time qualifiers. When the I64 linker produces an executable image, it copies the floating point controls from the object module that supplied the main entry point transfer address for the image into the image itself; this is called the “whole program floating-point mode” for the image. Then when the image is activated for execution, the hardware’s floating-point controls are

initialized according to this whole program floating point mode. It is expected that code that modifies the floating-point controls at run-time will be written to ensure that the whole program floating point mode settings get restored whenever control passes out of the section of code that required the specific setting of the controls at run-time.

When the `/IEEE_MODE` qualifier is applied to a compilation that does not contain a main program, the qualifier does have some effect: it can affect the evaluation of floating-point constant expressions, and it is used to set the `EXCEPTION_MODE` used by the math library for calls from that compilation.

The qualifier has no effect on the exceptional behavior of floating-point calculations generated as inline code for that compilation. Therefore, if floating point exceptional behavior is important to an application, all of its compilations, including the one containing the main program, should be compiled with the same `/FLOAT` and `/IEEE_MODE` settings.

Note that even on Alpha, setting `/IEEE_MODE=UNDERFLOW_TO_ZERO` requires the setting of a run-time status register; therefore, this setting needs to be specified on the compilation containing the main program in order to be effective in other compilations.

Finally, note that the `/ROUNDING_MODE` qualifier is affected in the same way as `/IEEE_MODE`, and is included in the whole program floating-point mode, and that because VAX floating point operations are actually performed using IEEE instructions, compilations that use VAX format floating-point exhibit the same whole program floating-point mode settings as compilations with `/IEEE_MODE=DENORM/ROUND=NEAREST`.

### 6.6.3 Predefined Macros

The compiler predefines a number of macros with the same meanings as in the native Alpha compiler, except that it does not predefine any of the macros that specify the Alpha architecture. Instead, it predefines the macros `__ia64` and `__ia64__`, as is the practice in the Intel and gcc compilers for I64. The change in floating-point representation from G\_FLOAT to IEEE is reflected in the macros that are predefined by default.

Some users have tried defining the macro `__ALPHA` explicitly by using `/DEFINE` or by putting it in a header file as a quick way to deal with source code conditionals that were written to assume that if `__ALPHA` is not defined, then the target must be a VAX system. Doing this causes the CRTL headers and other OpenVMS headers to take the wrong path for I64. You must not define any of the Alpha architecture predefined macros when using this compiler.

## 6.7 HP C++

HP C++ Version 7.1 is supported on OpenVMS I64. Refer to the HP C++ documentation for details and additional porting considerations.

### 6.7.1 Floating Point and Predefined Macros

Exactly the same considerations for floating-point defaults, `/IEEE_MODE` semantics, and predefined macros as described for the C compiler above apply to the C++ compiler.

## Preparing to Port Applications

### 6.7 HP C++

#### 6.7.2 Long Double

The long double type is always represented in 128-bit IEEE quad precision. The `/L_DOUBLE_SIZE=64` qualifier is ignored with a warning. Note that on Alpha, C++ library support for 64-bit long double was limited - code that requires a 64-bit floating point type should use double instead of long double.

#### 6.7.3 Object Model

The object model and name mangling schemes used by the C++ compiler on I64 differ significantly from those used on Alpha. The "ARM" object model is not available, the only object model is one that supports the ANSI/ISO C++ standard. However, this still differs from the `/MODEL=ANSI` object model implemented on Alpha, as the model on I64 is an implementation of the industry-standard I64 ABI. Programs that rely on the layout of C++ objects (non "POD" data), or on the external mangles names as encoded in the .obj file, will need to be reworked. Such programs are inherently highly implementation-dependent. But programs that use standard C++ features in a reasonable implementation-independent manner should not have difficulty in porting.

#### 6.7.4 Language Dialects

The "cfront" dialect is no longer supported (and will be removed from Alpha as well). Compilations specifying `/standard=cfront` will instead use the "relaxed\_ansi" dialect.

#### 6.7.5 Templates

On OpenVMS I64, .OBJ files are implemented in ELF format rather than EOBJ, and along with the I64 linker they support the notion of "COMDAT" sections that have been used for some time on both Windows and Unix platforms to resolve the issues of duplicate definitions at link-time that arise when using C++ templates and inline functions. On Alpha, these issues are handled by the repository mechanism, which arranges to present a single defining instance to the linker. On I64, no repository mechanism is needed to do this, as duplicates are discarded by the linker automatically. So the repository-based template instantiation options supported on Alpha are not supported on IPF. Alpha build procedures that attempt to manipulate objects in the repository will fail on I64 and will need to be changed (because there are no objects in the repository on I64, just the demangler database). In most cases, the reason for manipulating the repository directly in the build procedure has been eliminated by the compiler's use of COMDAT instantiation.

### 6.8 Java

Java™ Version 1.4.2-1 is supported on OpenVMS I64. There are no differences between Java on OpenVMS Alpha and OpenVMS I64.

### 6.9 Macro-32

Most VAX MACRO programs that compile on OpenVMS Alpha should recompile on OpenVMS I64 without modification. However, programs that call routines with nonstandard return values or programs that use the JSB instruction to call routines written in other languages must add some new directives to the source file. For more information, refer to the *HP OpenVMS MACRO Compiler Porting and User's Guide*.

## 6.10 HP Pascal

HP Pascal Version 5.9 is available for both Alpha and I64.

There are no known language differences between Pascal on OpenVMS Alpha and Pascal on OpenVMS I64. The only difference is that the default is for IEEE floating datatypes instead of VAX floating datatypes, which is true for all the HP I64 compilers.

### 6.10.1 /ARCH and /OPTIMIZE=TUNE Qualifiers

For the sake of “compile-and-go” compatibility, Alpha values for the /ARCH and /OPTIMIZE=TUNE qualifiers are accepted on the compiler invocation command. An informational message is printed saying they are ignored.

I64 values for /ARCH and /OPTIMIZE=TUNE are defined in the OpenVMS I64 compiler for development purposes only. Their behavior is unpredictable and they should not be used.



---

## Other Considerations

This chapter describes additional porting considerations.

### 7.1 Hardware Considerations

This section contains information about the Alpha, PA-RISC and the Itanium processor families. It also discusses processor-related development issues and provides examples using both C and assembly language for typical processor-dependent code.

#### 7.1.1 Intel Itanium Processor Family Overview

The Itanium architecture was developed by HP and Intel as a 64-bit processor for engineering workstations and e-commerce servers. The Itanium processor family uses a new architecture called Explicitly Parallel Instruction Computing (EPIC).

Some of the most important goals in the design of EPIC were instruction-level parallelism, software pipelining, speculation, predication, and large register files. These allow Itanium processors to execute multiple instructions simultaneously. EPIC exposes many of these parallel features to the compiler (hence "Explicitly" in the name). This increases the complexity of the compiler, but it allows the compiler to directly take advantage of these features.

EPIC is designed so future generations of processors can utilize different levels of parallelism, which increases the flexibility of future designs. This allows the machine code to express where the parallelism exists without forcing the processor to conform to previous machine widths.

For more information on assembly programming for the Intel Itanium processor family, see the *Intel® Itanium® Architecture Assembly Language Reference Guide*, available at:

[http://developer.intel.com/software/products/opensource/tools1/tol\\_whte2.htm](http://developer.intel.com/software/products/opensource/tools1/tol_whte2.htm)

#### 7.1.2 Alpha Processor Family Overview

Alpha is a 64-bit load/store RISC architecture that is designed with particular emphasis on the three elements that most affect performance: clock speed, multiple instruction issue, and multiple processors.

The Alpha architects examined and analyzed current and theoretical RISC architecture design elements and developed high-performance alternatives for the Alpha architecture.

The Alpha architecture is designed to avoid bias toward any particular operating system or programming language. Alpha supports the OpenVMS Alpha, Tru64 UNIX, and Linux operating systems and supports simple software migration for applications that run on those operating systems.

## Other Considerations

### 7.1 Hardware Considerations

Alpha was designed as a 64-bit architecture. All registers are 64 bits in length and all operations are performed between 64-bit registers. It is not a 32-bit architecture that was later expanded to 64 bits.

The instructions are very simple. All instructions are 32 bits in length. Memory operations are either loads or stores. All data manipulation is done between registers.

The Alpha architecture facilitates pipelining multiple instances of the same operations because there are no special registers and no condition codes.

The instructions interact with each other only by one instruction writing a register or memory and another instruction reading from the same place. That makes it particularly easy to build implementations that issue multiple instructions every CPU cycle.

Alpha makes it easy to maintain binary compatibility across multiple implementations and easy to maintain full speed on multiple-issue implementations. For example, there are no implementation-specific pipeline timing hazards, no load-delay slots, and no branch-delay slots.

The Alpha architecture reads and writes bytes between registers and memory with the LDBU and STB instructions. (Alpha also supports word read/writes with the LDWU and STW instructions.) Byte shifting and masking are performed with normal 64-bit register-to-register instructions, crafted to keep instruction sequences short.

As viewed from a second processor (including an I/O device), a sequence of reads and writes issued by one processor can be arbitrarily reordered by an implementation. This allows implementations to use multibank caches, bypassed write buffers, write merging, pipelined writes with retry on error, and so forth. If strict ordering between two accesses must be maintained, explicit memory barrier instructions can be inserted in the program.

The basic multiprocessor interlocking primitive is a RISC-style `load_locked`, `modify`, `store_conditional` sequence. If the sequence runs without interrupt, exception, or an interfering write from another processor, then the conditional store succeeds. Otherwise, the store fails and the program eventually must branch back and retry the sequence. This style of interlocking scales well with very fast caches and makes Alpha an especially attractive architecture for building multiple-processor systems.

A privileged architecture library (PALcode) is a set of subroutines that are specific to a particular Alpha operating system implementation. These subroutines provide operating-system primitives for context switching, interrupts, exceptions, and memory management. PALcode is similar to the BIOS libraries that are provided in personal computers.

For more information, refer to the *Alpha Architecture Handbook*, part number EC-QD2KC-TE.

### 7.2 Endianism Considerations

Users who are porting from OpenVMS Alpha or OpenVMS VAX to OpenVMS I64 do not need to be concerned with endianism issues; OpenVMS on all platforms uses little-endian for data storage and manipulation. This chapter is being made available more for reference to allow the reader to be aware of issues, in general, when porting between different hardware architectures where endianisms are different.




## Other Considerations

### 7.2 Endianism Considerations


Endianism refers to the way data is stored and defines how bytes are addressed in multibyte data types. This is important because if you try to read binary data on a machine that is of a different endianism than the machine that wrote the data, the results will be different. This is an especially significant problem when a database needs to be moved to a different endian architecture.

There are two types of endian machines: big-endian (forward byte ordering or most significant first, [MSF]) and little-endian (reverse byte ordering or least significant first, [LSF]). The mnemonics “big end in” and “little end in” can be useful when discussing endianism. Figure 7–1 and Figure 7–2 compare the two byte-ordering methodologies.

**Figure 7–1 Big-Endian Byte Ordering**



**Figure 7–2 Little-Endian Byte Ordering**



The Alpha microprocessor provides a little-endian environment for operating systems such as OpenVMS and Tru64 UNIX, while the PA-RISC microprocessor provides a big-endian environment for the HP-UX operating system.

On a little-endian operating system, such as OpenVMS, the little end, or least-significant byte is stored at the lowest address. This means a hexadecimal like 0x1234 is stored in memory as 0x34 0x12. The same is true for a 4-byte value; for example, 0x12345678 is stored as 0x78 0x56 0x34 0x12. A big-endian operating system does this in reverse fashion, so 0x1234 is stored as 0x12 0x34 in memory.

## Other Considerations

### 7.2 Endianism Considerations

Consider the number 1025 ( $2$  to the tenth power plus  $1$ ) stored in a 4-byte integer. This is represented in memory as follows:

```
little endian: 00000000 00000000 00000100 00000001b  
big endian: 00000001 00000100 00000000 00000000b
```

The Intel Itanium processor family architecture supports both big-endian and little-endian memory addressing. There are no problems reading and exchanging data between OpenVMS Alpha and OpenVMS I64 running on Itanium processors. If there is a need to exchange data between OpenVMS on an Alpha or I64 system and HP-UX on PA-RISC or the Itanium processor family, then byte swapping might be required.

---

## **Application Evaluation Checklist**

Use this checklist to guide you through the planning and porting process.

# Application Evaluation Checklist

## General Information

1. What does your application do (brief description)?

---

---

## Development History and Plans

2. a. Does the application currently run on other operating systems or hardware architectures? ☐ YES ☐ NO
- b. If yes, does the application currently run on the latest OpenVMS Alpha system? ☐ YES ☐ NO

[If your application already runs on multiple platforms, it is likely to be platform independent, and therefore, easier to port to OpenVMS I64. (However, if the application code is conditionalized per platform, then minor modifications might be required. For more information, see Section 4.8.1.) If you answer YES to b, your application will be easier to port to OpenVMS I64. If your application is running on OpenVMS VAX, you must first port it to OpenVMS Alpha. (For more information, refer to the *Migrating an Application from OpenVMS VAX to OpenVMS Alpha* manual.)]

3. When was the last time the application environment was completely recompiled and rebuilt from source? \_\_\_\_\_

- a. Is the application rebuilt regularly? ☐ YES ☐ NO
- b. How frequently? \_\_\_\_\_

[Applications that are not built frequently might require additional work before being ported. For example, changes in the development environment might cause incompatibilities. Prior to porting your applications, confirm that they can be built on the latest version of OpenVMS Alpha.]

4. Is the application actively maintained by developers who know it well? ☐ YES ☐ NO

List developer names and contact information:

Developer	Contact
_____	_____
_____	_____
_____	_____

5. a. How is a new build of the application tested or verified for proper operation?

---

---

- b. Do you have performance characterization tools to assist with optimization? ☐ YES ☐ NO

c. If yes, list the tools and version numbers used:

## Application Evaluation Checklist

Tool	Version
_____	_____
_____	_____
_____	_____

d. Which source-code configuration management tools are used?  
List the tools and their version numbers:

Tool	Version
_____	_____
_____	_____
_____	_____

6. Do you have a development and test environment separate from your production systems? ☐ YES ☐ NO
7. What procedures are in place for rolling in a new version of the application into production?

\_\_\_\_\_

\_\_\_\_\_

8. What are your plans for the application after porting it to OpenVMS I64?
- a. No further development ☐ YES ☐ NO
- b. Maintenance releases only ☐ YES ☐ NO
- c. Additional or changed functionality ☐ YES ☐ NO
- d. Maintain separate OpenVMS I64 and OpenVMS Alpha sources ☐ YES ☐ NO

[If you answer YES to a, you might want to consider translating the application. A YES response to b or c is reason to evaluate the benefits of recompiling and relinking your application, although translation is still possible. If you intend to maintain separate OpenVMS I64 and OpenVMS Alpha sources, as indicated by a YES to d, you might need to consider interoperability and consistency issues, especially if different versions of the application can access the same database.]

### Composition of the Application

Consider the size and components of your application that need to be ported.

9. How large is your application?
- How many modules? \_\_\_\_\_
- How many lines, disk blocks, or megabytes of code? \_\_\_\_\_
- How much disk space is required? \_\_\_\_\_
- [This will help you “size” the effort and the resources required for porting. Most helpful is knowing the number of disk blocks of code.]
10. a. Do you have access to all source files that make up your application? ☐ YES ☐ NO

## Application Evaluation Checklist

- b. If you are considering using HP Services, will it be possible to give HP access to these source files and build procedures? ☐ YES ☐ NO

[If you answer NO to a, translation may be your only porting option for the files with missing sources. Translate user-mode OpenVMS Alpha images to OpenVMS I64 images.

Note that if you have OpenVMS VAX images, first translate them to OpenVMS Alpha images. You cannot translate them directly to OpenVMS I64 images.

A YES answer to b allows you to take advantage of a greater range of HP porting services.]

11. a. List the languages used to write the application. (If multiple languages are used, give the percentages of each.)

Language	Percentage
_____	_____
_____	_____
_____	_____

- b. Do you use Macro-64, PL/I, Ada 83, or Fortran 77 on OpenVMS Alpha? ☐ YES ☐ NO

[If you use Macro-64 on OpenVMS Alpha, you must rewrite the code because no Macro-64 compiler exists for OpenVMS I64. If possible, HP recommends rewriting the code into a high-level language using documented system interfaces. For machine-specific code that must be written in assembly language, you must recode the Macro-64 code into OpenVMS I64 assembly code.

Similarly, PL/I, Fortran 77, and ADA 83 are not supported on OpenVMS I64. If your application has code written in PL/I, HP recommends rewriting it in another language such as C or C++. Update code written in Ada 83 to Ada 95, and code written in Fortran 77 to Fortran 90. Analyze your code to determine whether the newer compiler requires changes to the code.

In general, if the compilers are not available on OpenVMS I64, you must translate or rewrite code in a different language.

For information on availability of compilers and translators, see Chapters 5 and 6.]

12. Is there application documentation? ☐ YES ☐ NO

[If you answer yes, note that if any changes were required to the application, the documentation might have to be updated accordingly.]

### External Dependencies

Consider the system configuration, environment, and software required for developing, testing, and running your application.

13. What is the system configuration (CPUs, memory, disks) required to set up a development environment for the application?

---

---

## Application Evaluation Checklist

[This will help you plan for the resources needed for porting.]

14. What is the system configuration (CPUs, memory, disks) required to set up a typical user environment for the application, including installation verification procedures, regression tests, benchmarks, or work loads?

---

---

[This will help you determine whether your entire environment is available on OpenVMS I64.]

15. Does the application rely on any special hardware? ☐ YES ☐ NO

[This will help you determine whether the hardware is available on OpenVMS I64, and whether the application includes hardware-specific code.]

16. What version of OpenVMS Alpha does your application currently run on? \_\_\_\_\_

Does the application run on OpenVMS Alpha Version 7.3-2 or 7.3-1? ☐ YES ☐ NO

[HP recommends having your application running on the latest version of OpenVMS Alpha before porting it to OpenVMS I64.]

17. Does the application require layered and third-party products to run?

a. From HP (other than compiler RTLs): ☐ YES ☐ NO

List the HP layered products and versions:

HP Product	Version
_____	_____
_____	_____
_____	_____

b. From third parties: ☐ YES ☐ NO

List the third-party products and versions:

Third-Party Product	Version
_____	_____
_____	_____
_____	_____

[If you answer YES to a and are uncertain whether the HP layered products are yet available for OpenVMS I64, check with an HP support representative. For information about HP's commitment to support porting of software, see Section 3.4. If you answer YES to b, check with your third-party product supplier.]

18. a. Do you have regression tests for the application? ☐ YES ☐ NO

b. If yes, do they require any particular software product such as HP Digital Test Manager? ☐ YES ☐ NO

## Application Evaluation Checklist

[If you answer YES to a, you should consider porting those regression tests. If you answer YES to b, HP Digital Test Manager (part of the DECset product set) is available with this release of OpenVMS. If you require other tools for testing purposes, contact your HP support representative; see also Chapter 5. HP recommends a disciplined testing process to check regressions in product development.]

### Dependencies on the OpenVMS Alpha Architecture

Consider differences between OpenVMS Alpha and OpenVMS I64. The following questions will help you account for the most significant differences. Consider the changes that might be necessary to application components because of these differences. Any user-written code or assembly code that depends specifically on the OpenVMS Alpha architecture must be rewritten.

19. Does the application use OpenVMS VAX floating-point data types? ☐ YES ☐ NO

[If you answer YES, note that the default for OpenVMS I64 compilers is IEEE floating data types. OpenVMS I64 compilers provide for VAX floating support by automatically converting from OpenVMS VAX to IEEE floating data types. Slight precision differences might result. In addition, run-time performance will be somewhat slower (compared to the direct use of IEEE floating data types without conversion), depending on the extent to which the application is dominated by floating point calculations. For more information about floating point data types, see Section 4.8.4.]

20. Is the data naturally aligned? ☐ YES ☐ NO

[If your data is not naturally aligned, HP recommends that you align your data naturally to achieve optimal performance for data referencing. Unaligned data can degrade data referencing performance significantly. In certain instances, you must align data naturally. For more information about data structures and alignment, see Section 4.8.7 and refer to the *HP OpenVMS Programming Concepts Manual*.]

21. Does the application call OpenVMS system services that:

a. Modify the working set?

☐ YES ☐ NO

b. Change program control flow using SYS\$GOTO\_UNWIND?

☐ YES ☐ NO

[If you answer YES to a, make sure that calls to the SYS\$LKWSET and SYS\$LKWSET\_64 system services specify the required input parameters correctly.

If you answer YES to b, change any calls to SYS\$GOTO\_UNWIND to SYS\$GOTO\_UNWIND\_64. OpenVMS I64 does not support SYS\$GOTO\_UNWIND.

For more information, see Chapter 4.]

22. a. Does the application use multiple, cooperating processes? ☐ YES ☐ NO

If so:

b. How many processes?

\_\_\_\_\_



## Application Evaluation Checklist

c. What interprocess communication method is used? \_\_\_\_\_

- |                                   |                                    |                                |                                |
|-----------------------------------|------------------------------------|--------------------------------|--------------------------------|
| <input type="checkbox"/> \$CRMPSC | <input type="checkbox"/> Mailboxes | <input type="checkbox"/> SCS   | <input type="checkbox"/> Other |
| <input type="checkbox"/> DLM      | <input type="checkbox"/> SHM, IPC  | <input type="checkbox"/> SMG\$ | <input type="checkbox"/> STR\$ |

d. If you use global sections (\$CRMPSC) to share data with other processes, how is data access synchronized? \_\_\_\_\_

[This will help you determine whether you will need to use explicit synchronization, and the level of effort required to guarantee synchronization among the parts of your application. Use of a high-level synchronization method generally allows you to port an application most easily.]

23. Does the application currently run in a multiprocessor (SMP) environment? ☐ YES ☐ NO

[If you answer YES, it is likely that your application already uses adequate interprocess synchronization methods.]

24. Does the application use AST (asynchronous system trap) mechanisms? ☐ YES ☐ NO

[If you answer YES, you should determine whether the AST and main process share access to data in process space. If so, you may need to explicitly synchronize such accesses. Note that this is equally true for applications on Alpha. For more information about ASTs, see *HP OpenVMS Programming Concepts Manual*.]

25. a. Does the application contain condition handlers? ☐ YES ☐ NO  
 b. Do any handlers modify data in the mechanism array? ☐ YES ☐ NO

[The OpenVMS Alpha and I64 mechanism array formats differ significantly. Be sure to modify any applications that modify register values in the mechanism array. For more information, see the *HP OpenVMS Guide to Upgrading Privileged-Code Applications*.]

26. a. Does the application run in privileged mode or link against SYS\$BASE\_IMAGE? ☐ YES ☐ NO

If so, why? \_\_\_\_\_

- b. Does the application depend on OpenVMS internal data structures or interfaces? ☐ YES ☐ NO

[Applications that link against the OpenVMS executive or run in privileged mode might require additional porting work. Undocumented interfaces in SYS\$BASE\_IMAGE might have changed on OpenVMS I64.

Applications that depend on OpenVMS internal data structure definitions (as defined in SYS\$LIBRARY:LIB.INCLUDE, SYS\$LIBRARY:LIB.L32, SYS\$LIBRARY:LIB.MLB, SYS\$LIBRARY:LIB.R64, SYS\$LIBRARY:LIB.REQ, and SYS\$LIBRARY:SYS\$LIB\_C.TLB) might also require additional porting work, as some internal data structures have changed on OpenVMS I64.]

## Application Evaluation Checklist

27. Does the application use connect-to-interrupt mechanisms? ☐ YES ☐ NO  
If yes, with what functionality? \_\_\_\_\_  
[Connect-to-interrupt is not supported on OpenVMS I64 systems. Contact an HP representative if you need this feature.]
28. Does the application create or modify machine instructions? ☐ YES ☐ NO  
[Guaranteeing correct execution of instructions written to the instruction stream requires great care on OpenVMS I64. Any code dealing with specific machine instructions must be rewritten.]
29. Does the application include any other user-written code or assembler code that depends specifically on the OpenVMS Alpha architecture? ☐ YES ☐ NO  
[If you answer YES, rewrite such code. Optimally, rewrite it so that it is independent of architecture.]
30. Does the application include or depend on any conditionalized statements in source code or build files, or any code that includes logic that assumes it is running on either a VAX or an OpenVMS Alpha system? ☐ YES ☐ NO  
[If you answer YES, you might have to modify the conditional directives slightly. For example, change “IF ALPHA” to “IF ALPHA OR IPF”, or change “IF NOT ALPHA” to “IF VAX”. For more information, see Section 4.8.1.]
31. What parts of the application are most sensitive to performance? I/O, floating point, memory, realtime (that is, interrupt latency, and so on).  
\_\_\_\_\_  
\_\_\_\_\_  
[This will help you determine how to prioritize work on the various parts of your application and allow HP to plan performance enhancements that are most meaningful to customers.]
32. Does the application use any OpenVMS Alpha Calling Standard routines? ☐ YES ☐ NO  
[If you answer YES, note that applications that use the calling standard data structure named LIBICB, or that make use of the calling standard routines such as LIB\$GET\_CURRENT\_INVO\_CONTEXT, LIB\$GET\_PREVIOUS\_INVO\_CONTEXT, LIB\$PUT\_INVO\_CONTEXT, and LIB\$GET\_INVO\_HANDLE, will have to be modified. These routines have been renamed in OpenVMS I64 due to changes in the LIBICB data structure. In addition, the calling standard has changed. Applications with direct knowledge of the OpenVMS Alpha calling standard must be modified for OpenVMS I64. For more information on the differences between OpenVMS Alpha and I64 calling standards, see Section 2.1.1. For more information about OpenVMS calling standards, see the *HP OpenVMS Calling Standard*.]

## Application Evaluation Checklist

33. Does the application use non-standard routine linkage declarations? ☐ YES ☐ NO

[If you answer YES, note that by definition applications with non-standard linkages know the OpenVMS Alpha calling standard and might need modification to conform to the OpenVMS I64 calling standard.]

34. a. Does the application depend on the format or content of an object file? ☐ YES ☐ NO  
b. Does the application depend on the format or content of an executable image? ☐ YES ☐ NO  
c. Does the application depend on the debug symbol table content of an image? ☐ YES ☐ NO

[If you answer YES to a, b, or c, note that applications that depend on any of these image and object file data structures must be modified. The industry-standard object and image file layout called ELF has been adopted by OpenVMS I64. The industry-standard DWARF debug symbol table format has also been adopted. Applications with dependencies on OpenVMS Alpha's object, image or debug symbol table formats must be modified. For more information, see Section 4.8.3.4.]

35. Does the application use the C *asm* statement? ☐ YES ☐ NO

[If you answer YES, note that applications that use the C *asm* statement have Alpha instruction dependencies and must be modified. The HP C compiler for OpenVMS I64 has added many new builtin functions that minimize the need for using *asm* statements.]

36. Does the application use BLISS register Builtins? ☐ YES ☐ NO

[If you answer YES, note that applications that use the BLISS BUILTIN statement to declare register builtins must be modified.]

### Porting Your Application

Once the preliminary evaluation and planning stages are completed, these are the main porting tasks.

37. Have you upgraded your OpenVMS Alpha system to the latest version available? ☐ YES ☐ NO

[If you answer NO, first upgrade your OpenVMS Alpha system before beginning to port the application.]

38. a. Have you test compiled your application on the latest version of OpenVMS Alpha, using the most recent version of the compiler? ☐ YES ☐ NO  
b. Have you corrected any problems detected during the compile? ☐ YES ☐ NO

## Application Evaluation Checklist

[If you answer NO to a, compile your application on the latest OpenVMS Alpha system, using the latest compiler available. If you answer YES to a but have rewritten some code since, compile again. After compiling, correct any problems detected. When you have finished this phase, and when you have responded to all other considerations detailed thus far in this checklist and in Chapters 4 through 7, your application is ready for porting to OpenVMS I64.]

39. Copy the application source modules and all related code to the OpenVMS I64 system. Have you copied the modules and related code to OpenVMS I64? ☐ YES ☐ NO
40. Compile, link, and run the application on the OpenVMS I64 system.  
Application is compiled, linked, and running? ☐ YES ☐ NO  
[For details, see Chapter 5.]
41. Perform unit testing on the OpenVMS I64 system.  
Have you completed the unit testing? ☐ YES ☐ NO  
[For details, see Chapter 5.]
42. Perform system testing on the OpenVMS I64 system.  
Have you completed the system testing? ☐ YES ☐ NO  
[For details, see Chapter 5.]
43. Perform regression testing on the OpenVMS I64 system.  
Have you completed the regression testing? ☐ YES ☐ NO  
[For details, see Chapter 5.]
44. Did testing succeed? ☐ YES ☐ NO
- [If you answer NO, resolve the problems, then recompile, relink, and retest the application. For details, see Chapter 5.]

---

## Unsupported Layered Products

HP has not plans to port some layered products to OpenVMS I64. Table B–1 lists these layered products and the replacement products recommended by HP.

**Table B–1 Layered Products Not Ported to OpenVMS I64**

Product	Suggested Replacement
BASEstar Classic	BASEstar Open
HP Ada Compiler for OpenVMS Alpha Systems	GNAT Ada for OpenVMS Integrity Servers—from Ada Core Technologies
Pathworks 32	Advanced Server



---

## Porting Application-Specific Stack-Switching Code to I64

Many applications support small, private threading packages that switch stacks to perform multiple tasks within the context of a single process. Many of these packages use one or more small routines written in assembly language (Macro-32 on VAX or Macro-64 on Alpha) to manipulate the stack pointers and to cause, in effect, a context switch within a single process. Typically, the ability to stall and restart the tasks is also required. The replaced stack can be in any mode appropriate to the required task.

On I64, switching stacks is much more difficult. The I64 architecture includes two stacks, a memory stack and a register stack engine backing store (commonly called the *RSE stack* or simply the *register stack*.) The RSE stack is maintained by hardware and the architecture provides support for asynchronous manipulation of the stack. Also, the I64 architecture includes many specialized hardware registers (called control and application registers) that must be maintained across a context switch. In addition, the IAS assembler is not a supported part of the OpenVMS distribution.

To accommodate these types of applications, OpenVMS offers a set of system routines called KP services. Originally written to support device drivers written in higher-level languages, the KP services have been expanded to work in any mode and at any IPL. While the KP model may not match the needs of all private stack-switching code, HP strongly suggests moving to this model.

KP services are available on both Alpha and I64, allowing for common code. There are some differences in implementation; however, whenever possible, options valid on only one architecture are ignored on the other.

### C.1 Overview of KP Services

The KP model is most accurately described as a co-routine model with voluntary stall and resumption. A code stream running in process or system context desires to start another stream (called the *KP routine*) in the same mode but without disturbing the existing stack context and allowing the KP routine to stall and resume as necessary while maintaining its own distinct context. Resumption of a KP routine can be asynchronous to the original code stream that started the KP routine.

The easiest view of a KP routine is that of a routine that operates on a private stack (or on I64, a pair of stacks.) The routine can be stalled and resumed. To a compiler, the routines to start, stall, and resume simply appear as outbound procedure calls which obey the OpenVMS calling standard. The saving of state and stack switch takes place entirely within the context of the KP routines. When a code stream gives up control, the resumption point is always as if the call that caused transfer of control completed.

## Porting Application-Specific Stack-Switching Code to I64

### C.1 Overview of KP Services

The base support includes routines to assist in the allocation of necessary stacks and data structures and routines to start, stall, restart, and terminate a KP routine.

A KP routine can start another KP routine.

The basic KP support routines are:

- `EXE$KP_START`—Start a KP routine.
- `EXE$KP_STALL_GENERAL`—stall the current routine and return to the last caller of `EXE$KP_START` or `EXE$KP_RESTART`.
- `EXE$KP_RESTART`—Restart a KP routine that has been stalled by a call to `EXE$KP_STALL_GENERAL`.
- `EXE$KP_END`—Terminate the KP routine.

The routines and required data structures are described in Section C.2.2.

#### C.1.1 Terminology

The **memory stack** is the stack that is pointed to the stack pointer register. The memory stack is the only stack on Alpha systems.

The **register stack** is the informal name for the **register stack engine (RSE) backing store**. The I64 architecture includes 96 stacked registers that are used for argument passing between routines and that can be used for local storage within a routine as well. The hardware maintains the stacked register state and flushes registers to the backing store as needed.

A **KP routine** is a sequence of code that runs on a private stack or stacks. It is started via a call to `EXE$KP_START` and is terminated by an implicit or explicit call to `EXE$KP_END`.

A **KPB** is the data structure that describes the stack or stacks and that maintains state and context information for the KP routine.

A KPB is **valid** if it has been used to activate a KP routine via `EXE$KP_START`. `EXE$KP_END` marks the KPB as **invalid**. The initial state of a KPB is invalid.

A KPB is **active** when the KP routine has been started via `EXE$KP_START` or resumed via `EXE$KP_RESTART`. `EXE$KP_STALL_GENERAL` and `EXE$KP_END` mark the KPB as **inactive**. The initial state of a KPB is inactive.

#### C.1.2 Stacks and Data Structures

On I64, three pieces of memory are associated with a KP routine—a memory stack, an RSE stack, and a KPB, which is the data structure that ties them all together. On Alpha, there is no RSE stack and all parameters and fields related to the RSE stack are ignored.

The KPB and stacks each must be allocated from an appropriate region, and at an appropriate mode, protection, and ownership to match the mode in which the KP routine will execute. When porting an existing application, it is expected that the application already allocates an appropriate memory stack. The existing memory stack allocation routine can be adapted to the KP API. As with previous architectures, the memory stack is accessed from the highest address (the base of the stack) to the lowest address.



# Porting Application-Specific Stack-Switching Code to I64

## C.1 Overview of KP Services

RSE stacks are typically allocated from 64-bit space because the register stack engine is a new entity with no previous 32-bit dependencies. A number of allocation routines have been supplied that should cover most common application needs. The RSE stack is accessed from the lowest address to the highest address.

Table C–1 offers guidelines for allocation by mode and scope of the application.

**Table C–1 Allocation Guidelines by Mode and Scope**

Mode-Scope <sup>1</sup>	KPB	Memory Stack	Register Stack
Kernel—System EXE\$KP_ALLOC_ KPB <sup>2</sup>	Nonpaged pool KW EXE\$ALONONPAGED	S0/S1 KW EXE\$KP_ALLOC_ MEM_STACK	S2 KW EXE\$KP_ALLOC_ RSE_STACK <sup>3</sup>
Kernel—Process	Nonpaged pool or P1 KW EXE\$ALONONPAGED or EXE\$ALOP1PROC	P1—Permanent KW \$CREATE_ REGION/\$CRETVA	P2—Permanent KW EXE\$KP_ALLOC_ RSE_STACK_P2
Kernel—Image	P1 KW EXE\$ALOP1IMAG	P1—Nonpermanent KW \$CREATE_ REGION/\$CRETVA	P2—Nonpermanent KW \$CREATE_ REGION/\$CRETVA
Exec—Process	P1 EW EXE\$ALOP1PROC	P1—Permanent EW \$CREATE_ REGION/\$CRETVA	P2—Permanent EW EXE\$KP_ALLOC_ RSE_STACK_P2
Exec—Image	P1 EW EXE\$ALOP1IMAG	P1—Nonpermanent EW \$CREATE_ REGION/\$CRETVA	P2—Nonpermanent EW \$CREATE_ REGION/\$CRETVA
Super—Process	P1 SW EXE\$ALOP1PROC	P1—Permanent SW \$CREATE_ REGION/\$CRETVA	P2—Permanent SW EXE\$KP_ALLOC_ RSE_STACK_P2
Super—Image	P1 SW EXE\$ALOP1IMAG	P1—Nonpermanent SW \$CREATE_ REGION/\$CRETVA	P2—Nonpermanent SW \$CREATE_ REGION/\$CRETVA
User—Image	P0 UW Heap/Malloc/LIB\$GET_ VM	P0—Nonpermanent <sup>4</sup> UW EXE\$KP_ALLOC_ MEM_STACK_USER	P2—Nonpermanent UW EXE\$KP_ALLOC_ RSE_STACK_P2

<sup>1</sup> Image scope terminates at image exit. Process scope terminates at process exit and will survive image rundown. System scope does not require process context.

<sup>2</sup> EXE\$KP\_ALLOC\_KPB allocates kernel mode KPB and kernel mode stacks in a single call.

<sup>3</sup> EXE\$KP\_ALLOC\_RSE\_STACK\_P2 creates permanent regions.

<sup>4</sup> Note that permanent memory regions may not be created in user mode.

### C.1.3 KPBs

The KPB is a data structure that is used to maintain the necessary context between the initiating code stream and the KP routine. The KPB is semitransparent. Some fields are maintained by the application, some by the KP routines and some are shared. The KP routines assume the KPB was zeroed on allocation and, thus that any nonzero field contains legitimate data.

## Porting Application-Specific Stack-Switching Code to I64

### C.1 Overview of KP Services

The structure definitions for a KPB are defined by the `$KPBDEF` macro for Macro-32 and `KPBDEF.H` for C. The KPB definitions are considered system-internal and thus supplied in `LIB.MLB` and `SYS$LIB_C.TLB`. For BLISS, `LIB.REQ` or `LIB.L32/LIB.L64` contain the KPB definitions.

The KPB is a variable-length structure consisting of a number of areas or substructures. Not all areas are required. The areas are:

- Base area
- Scheduling area
- VEST area
- Spinlock area
- Debug area
- User parameter area

The base area is required. It contains a standard structure header, the stack sizes and base addresses, flags (including what other areas are present), the memory stack pointer for the nonactive code stream, pointers to the other areas, and additional fields required by the base KP routines.

The scheduling area includes pointers to stall, restart and end handling routines, a fork block and a pointer to an additional fork block. With the exception of the end routine, most are required by high-IPL driver-level code only. Callers of `EXE$KP_USER_ALLOC_KPB` must supply an end routine to perform necessary cleanup of the allocated memory.

The VEST and spinlock areas are used primarily by the driver code.

The debug area offers limited tracing capability implemented in the driver support routines.

The user parameter area is simply undefined storage allocated contiguously to the other areas. The application is free to use this memory for its own needs.

#### C.1.4 Supplied KPB Allocation Routines

The operating system supplies two standard allocation routines for KPBs with associated stacks. The original kernel-mode, driver level interface has been retained unchanged from Alpha so that device drivers using the KP interface do not require source changes in this area. In addition, a mode-independent routine is available. The mode-independent routine calls application-specified routines to allocate the KPB and each of the stacks. Most new applications and applications porting to the KP API will use the latter routine.

Both the kernel and mode-independent routines initialize the KPB. C prototypes for all the supplied routines can be found in the head file `EXE_ROUTINES.H`.

#### C.1.5 Kernel Mode Allocation

The format for kernel mode allocation is as follows.

`EXE$KP_ALLOCATE_KPB kpb, stack_size, flags, param_size`

## Porting Application-Specific Stack-Switching Code to I64

### C.1 Overview of KP Services

#### C prototype

```
status = EXE$KP_ALLOCATE_KPB( KPB_PPS kpb,  
    int stack_size,  
    int flags,  
    int param_size)
```

For kernel mode use only. This routine has the same prototype as the original Alpha routine.

On I64, RSE stack size = memory stack size.

---

#### Note

The stack size is in bytes.

---

#### Parameters

- **KPB**—Address of a longword to receive the address of the allocated data structure. By convention within OpenVMS APIs, 32-bit pointers are passed by 32-bit address (short pointer to short pointer to struct KPB).
- **STACK\_SIZE**—A 32-bit value denoting the size of the stacks to be allocated in bytes. Passed by value. The supplied value is rounded up to a multiple of the hardware page size. In any case, the minimum number of pages allocated will be no smaller than the SYSGEN parameter KSTACKPAGES. The stack is allocated on page boundaries with unmapped guard pages on both ends. The memory stack is allocated in 32-bit S0/S1 address space.  
On I64 systems, the RSE stack is sized identically to the memory stack and is allocated from 64-bit S2 address space.
- **FLAGS**—Longword bitmask of flags. Passed by value. Table C–2 describes the flags that are defined for this parameter.

**Table C–2 Kernel Mode Allocation Flags**

Flag	Description
KP\$M_VEST	OpenVMS system KPB. In general, this flag should be set.
KP\$M_SPLOCK	Allocate a spinlock area within the KPB.
KP\$M_DEBUG	Allocate a debug area within the KPB.
KP\$M_DEALLOC_AT_END	KPB should be automatically deallocated when the kernel process routine terminates.
KP\$M_SAVE_FP (IA64 only)	Save floating-point context as well as the general registers. Certain operations, such as integer multiplication and division on IA64, can be implemented using floating-point operations. These operations use the minimal floating-point register set, which is by definition not part of the preserved register set. If the application uses only the minimal floating-point register set, this bit need not be set. If the application uses floating-point data, this bit must be set to preserve the correct floating point context across the stack switch.

(continued on next page)

## Porting Application-Specific Stack-Switching Code to I64

### C.1 Overview of KP Services

**Table C–2 (Cont.) Kernel Mode Allocation Flags**

Flag	Description
KP\$M_SET_STACK_LIMITS	Call \$SETSTK_64 at every stack switch. Process-scope applications should always set this flag because condition handling requires accurate stack limit values.

- **PARAM\_SIZE**—Longword by value. Size, in bytes, of the user parameter area in the KPB. Pass zero if no parameter area is required.

Return value (status):

```
SS$ _NORMAL
SS$ _INSMEM
SS$ _INSFARG
SS$ _INSFRPGS
```

#### C.1.6 Mode-Independent Allocation

The syntax for mode-independent allocation is as follows.

```
EXE$KP_USER_ALLOC_KPB kpb, flags, param_size, *kpb_alloc, mem_stack_bytes,
*memstk_alloc, rse_stack_bytes, *rsestk_alloc, *end_rtn
```

##### C prototype

```
status = EXE$KP_USER_ALLOC_KPB( KPB_PPS kpb, int flags,
    int param_size,
    int (*kpb_alloc)(),
    int mem_stack_bytes,
    int (*memstk_alloc)(),
    int rse_stack_bytes,
    int (*rsestk_alloc)(),
    void(*end_rtn)())
```

##### Parameters

- **KPB**—Address of a longword to receive the address of the allocated data structure. By convention within OpenVMS APIs, 32-bit pointers are passed by 32-bit address (short pointer to short pointer to struct KPB).
- **FLAGS**—Longword bitmask of flags. Passed by value. Table C–3 describes the flags that are defined for this parameter.

**Table C–3 Mode-Independent Allocation Flags**

Flag	Description
KP\$M_VEST	OpenVMS system KPB. In general, this flag should be set.
KP\$M_SPLOCK	Allocate a spinlock area within the KPB.
KP\$M_DEBUG	Allocate a debug area within the KPB.
KP\$M_DEALLOC_AT_END	KPB should be automatically deallocated when the kernel process routine terminates.

(continued on next page)

## Porting Application-Specific Stack-Switching Code to I64

### C.1 Overview of KP Services

**Table C–3 (Cont.) Mode-Independent Allocation Flags**

Flag	Description
KP\$M_SAVE_FP	(I64 only; ignored on Alpha) Save floating-point context as well as the general registers. Certain operations such as integer multiplication and division on I64 systems can be implemented using floating-point operations. These operations use the minimal floating-point register set, which is by definition not part of the preserved register set. If the application uses only the minimal floating-point register set, this bit need not be set. If the application uses floating-point data, this bit must be set to preserve the correct floating-point context across the stack switch.
KP\$M_SET_STACK_LIMITS	Call \$SETSTK_64 at every stack switch. Process-scope applications should always set this flag because condition handling requires accurate stack limit values.

- **PARAM\_SIZE** – Longword by value. Size, in bytes, of the user parameter area in the KPB. Pass zero if no parameter area is required.
- **KPB\_ALLOC** – Address of a procedure descriptor for the routine to allocate the KPB. The allocation routine will be called with two parameters, a length and a longword to receive the address of the allocated KPB. The length parameter is in bytes and is passed by reference. The allocation routine must allocate at least the required number of bytes out of 32-bit space appropriate for the mode in which the KP routine will run. The address is expected to be at least quadword aligned and the actual allocation size must be written back into the length argument.
- **MEM\_STACK\_BYTES** – 32-bit value denoting the size of the memory stack to be allocated in bytes. Passed by value. The supplied value is rounded up to a multiple of the hardware page size.
- **MEMSTK\_ALLOC** – Address of a procedure descriptor for a routine to allocate the memory stack. Section C.1.7 describes the format and required actions of this routine.
- **RSE\_STACK\_BYTES** – 32-bit value that denotes the size, in bytes, of the memory stack to be allocated. Passed by value. The supplied value is rounded up to a multiple of the hardware page size.
- **RSESTK\_ALLOC** – Address of a procedure descriptor for a routine to allocate the memory stack. Section C.1.7 describes the format and required actions of this routine.
- **END\_RTN** – Address of a procedure descriptor for the end routine. The end routine is called when the KP routine terminates, either by calling EXE\$KP\_END or by returning. An end routine is required to either cache or deallocate the stacks and KPB.

## Porting Application-Specific Stack-Switching Code to I64

### C.1 Overview of KP Services

#### C.1.7 Stack Allocation APIs

The stack allocation routines have identical APIs for both memory and RSE stack allocation. The routine is called with a 64-bit address of the allocated KPB and an integral number of hardware-specific pages (not pagelets) to allocate.

The syntax for specifying stack allocation routines is as follows:

status = *alloc-routine* (KPB\_PQ kpb, const int stack\_pages)

- **KPB**—64-bit address of a previously allocated KPB. Passed by 64-bit reference.
- **STACK\_PAGES**—Integral number of whole pages to allocate. Passed by 32-bit value.

The allocation routine is expected to allocate page-aligned address space. While not strictly necessary, it is strongly suggested that the stack be protected by no-access guard pages on both ends. HP also recommends that the minimum stack size be at least the value of the SYSGEN parameter KSTACKPAGES (global cell SGN\$GL\_KSTACKPAG). This allows a certain measure of control of the stack size without necessitating recompilation of the application. Also, stack usage on I64 is significantly different than on previous architectures, and the previously allocated stack size might not be adequate.

The memory stack allocation routine must set the following KPB fields as follows:

- **KPB\$IS\_STACK\_SIZE**—The size of the stack in bytes, not including guard pages.
- **KPB\$PQ\_STACK\_BASE**—The address of the stack base. For the memory stack, this is the address of the byte past the end of the allocated stack. In other words, this is the sum of the allocated address plus the size of the allocation in bytes.

The memory stack allocation routine can set the following KPB field as follows:

- **KPB\$Q\_MEM\_REGION\_ID**—The region ID returned by the \$CREATE\_REGION system service. This information is required to deallocate the stack.

The RSE stack allocation routine must set the following KPB fields as follows:

- **KPB\$IS\_STACK\_SIZE**—The size of the stack in bytes, not including guard pages.
- **KPB\$PQ\_STACK\_BASE**—The address of the stack base. For the RSE stack, this is the lowest allocated address.

The RSE stack allocation routine can set the following KPB field as follows:

- **KPB\$Q\_REGION\_ID**—The region ID returned by the \$CREATE\_REGION system service. This information is required to deallocate the stack.

Both routines must return status to the calling routine.

#### C.1.8 System-Supplied Allocation and Deallocation Routines

The system supplies a number of standard allocation routines. The routines adhere to the KP allocation API and can be used in place of a user-written routine if they meet the needs of the application. The following allocation routines are available:

- **EXE\$KP\_ALLOC\_MEM\_STACK** (kernel mode, S0/S1 space)
- **EXE\$KP\_ALLOC\_MEM\_STACK\_USER** (user mode, P0 space)

## Porting Application-Specific Stack-Switching Code to I64

### C.1 Overview of KP Services

- EXE\$KP\_ALLOC\_RSE\_STACK (caller's mode, S2 space)
- EXE\$KP\_ALLOC\_RSE\_STACK\_P2 (caller's mode, P2 space)

The following deallocation routines are supplied. All deallocation routines take a single argument, the address of the KPB.

- EXE\$KP\_DEALLOCATE\_KPB
  - Only KPBs allocated by EXE\$KP\_ALLOCATE\_KPB
  - Deallocates stacks and KPB
- EXE\$KP\_DEALLOC\_MEM\_STACK
- EXE\$KP\_DEALLOC\_MEM\_STACK\_USER
- EXE\$KP\_DEALLOC\_RSE\_STACK
- EXE\$KP\_DEALLOC\_RSE\_STACK\_P2

#### C.1.9 End Routine

The end routine is called by the KP services when EXE\$KP\_END is called, either explicitly or by reaching the end of the KP routine. Since EXE\$KP\_USER\_ALLOC\_KPB allows the specification of arbitrary allocation routines, the end routine must either cache the KPB for future use by the application or call the necessary deallocation routines for the stacks and KPB.

The end routine is called with two parameters as follows:

void *end\_routine* (KPB\_PQ KPB, int status)

- **KPB**—64-bit address of a previously allocated KPB. Passed by 64-bit reference.
- **STATUS**—optional status value. If EXE\$KP\_END was called explicitly, the status value is that supplied as the second argument to EXE\$KP\_END. If the second argument is omitted, the value SS\$NORMAL is supplied. If the KP routine terminated by returning, the value supplied is the return value of the KP routine.

## C.2 KP Control Routines

Once the KPB and stacks are allocated, the four routines that determine the state of a KP routine can be used.

### C.2.1 Overview

A KP routine is initiated by a call to EXE\$KP\_START. While running, the KP routine may elect to give up control by calling EXE\$KP\_STALL\_GENERAL. A stalled KP routine can be resumed by calling EXE\$KP\_RESTART. The KP routine is terminated either by an explicit call to EXE\$KP\_END or by returning from the routine. In the latter case, the KP services call EXE\$KP\_END implicitly.

When a KP routine starts, the current thread of execution state is saved onto the current stack, the KP stack is loaded, and the KP routine is called.

When the KP routine stalls, the KP routine's context is saved onto the KP stack, the stack is switched to the original stack, and the main routine's context is loaded from stack. The result is that a stall causes the original routine to return from the last call to EXE\$KP\_START or to EXE\$KP\_RESTART.

## Porting Application-Specific Stack-Switching Code to I64

### C.2 KP Control Routines

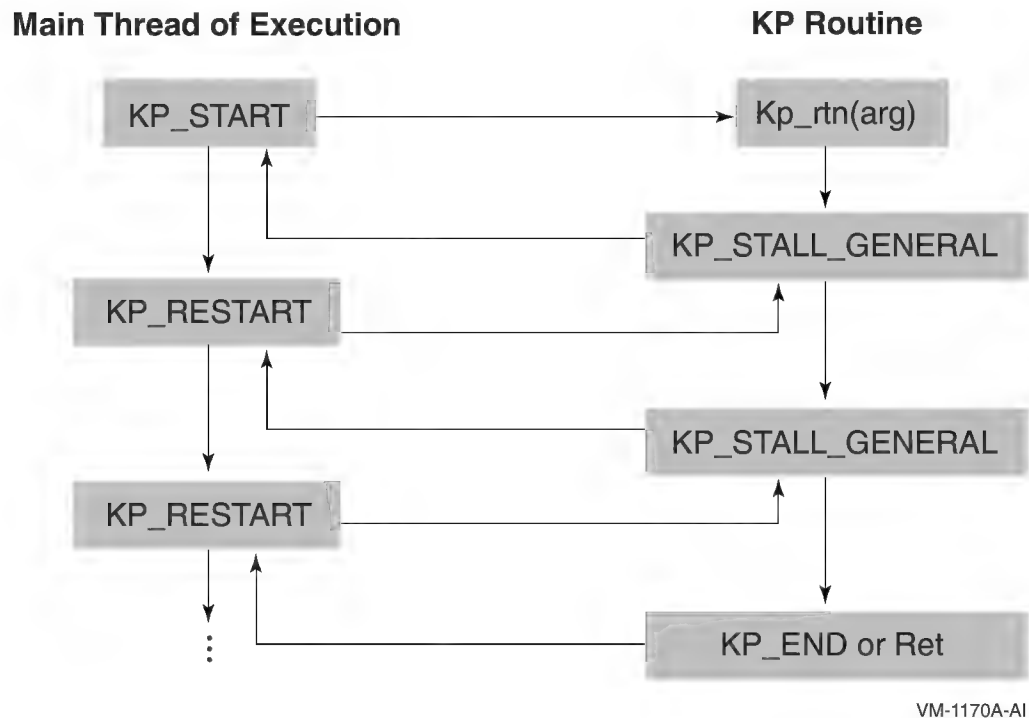
When the KP routine is restarted, the current context is saved onto the current stack, the stack is switched to the KP routine's stack and the KP routine's context is restored from the stack. The KP routine returns from the last call to EXE\$KP\_STALL\_GENERAL.

The stall/resume sequence can occur zero or more times. The KP routine is not required ever to stall. A stalled routine cannot stall again until it has been resumed. A running KP routine cannot be restarted until it has stalled. The full checking version of SYSTEM\_PRIMITIVES.EXE enforces these rules. Failure to follow these rules can result in a KP\_INCONSTATE bugcheck, depending on the mode in which the KP routine is running.

When the KP routine terminates, either by explicitly calling EXE\$KP\_END or by returning to the caller, no current context is saved, the stack is switched, and the original thread context is restored. At this point, if the DEALLOCATE\_AT\_END flag is set (kernel mode only) or if an end routine address has been supplied, the appropriate action takes place. The original thread returns from the call that started or restarted the KP routine.

Figure C-1 shows the overall code flow.

**Figure C-1 KP Routine Execution**



#### Note

While the main thread of execution is depicted as a continuous stream in the above diagram, the actual thread of execution may include asynchronous components. The application is required to perform all



necessary synchronization as well as maintaining the necessary scope of the KPB.

---

## C.2.2 Routine Descriptions

This section describes the routines.

### C.2.2.1 EXE\$KP\_START

*Syntax:*

status = EXE\$KP\_START(kpb, routine, reg-mask)

- **kpb**—Address of a previously allocated and initialize KPB. Passed by 32-bit reference.
- **routine**—KP routine address
- **reg-mask**—Register save mask. Longword. Passed by value. This argument is read on Alpha only. The I64 interface supports only the OpenVMS Calling Standard for register preservation. The constant KPREG\$K\_HLL\_REG\_MASK, defined in KPBDEF, can be used for calls from higher-level languages on Alpha to specify the calling standard register mask.

This routine suspends the current thread of execution, swaps to the new stacks; and calls the specified routine. The KP routine is called with a single argument—the 32-bit address of the supplied KPB. The KPB must be invalid and inactive.

### C.2.2.2 EXE\$KP\_STALL\_GENERAL

*Syntax:*

status = EXE\$KP\_STALL\_GENERAL(kpb)

- **kpb**—Address of the KPB passed at the start of this routine. Passed by 32-bit reference.

This routine stalls the current thread of execution, saving context onto the KP stack, and returns to the most recent call that started or restarted this routine. The KPB must be valid and active.

The return status from this routine is supplied by the routine that restarts this procedure.

### C.2.2.3 EXE\$KP\_RESTART

*Syntax:*

EXE\$KP\_RESTART(kpb [, thread\_status])

- **kpb**— Address of the KPB last used to stall this routine. Passed by 32-bit reference.
- **thread\_status** – Status value to be supplied as the return value for the call to EXE\$KP\_STALL\_GENERAL that last stalled the KP routine. Passed by value. Optional. If this parameter is omitted, SS\$\_NORMAL is returned.

## Porting Application-Specific Stack-Switching Code to I64

### C.2 KP Control Routines

This routine causes the stalled KP routine to restart by returning from the last call to EXE\$KP\_STALL\_GENERAL. Note that this may be a completely asynchronous operation with respect to the original thread of execution that started the KP routine. The KPB must be valid and inactive.

#### C.2.2.4 EXE\$KP\_RESTART

*Syntax:*

status = EXE\$KP\_END(kpb [, status])

- **kpb**—Address of the KPB last used to start or restart this routine. Passed by 32-bit reference.
- **status**—Status value to be supplied to the end routine, if any was specified in the KPB. Passed by value. Optional. If this parameter is omitted, SS\$NORMAL is returned.

This routine terminates the KP routine, returning control to the last thread of execution that started or restarted the KP routine. The KPB must be valid and active. The KPB is marked as invalid and inactive and cannot be used in subsequent calls to EXE\$KP\_RESTART or EXE\$KP\_STALL\_GENERAL without first calling EXE\$KP\_START to start another KP routine.

Instead of calling EXE\$KP\_END, the KP routine can return to the caller. Returning to the caller causes the KP code to call KP\_END automatically. In that case, the return status from the KP procedure is used for the status argument.

## C.3 Design Considerations

- KP routines run in a single process. The KP services by themselves offer no parallelism benefits on multiple-CPU systems.
- All calls to EXE\$KP\_STALL\_GENERAL, EXE\$KP\_RESTART and EXE\$KP\_END must occur in the same mode as the call to EXE\$KP\_START. Mode changes can occur between calls, but all calls must be made at the same mode.
- Multiple KPBs can be valid at the same time.
- Multiple KPBs can be active at the same time. This implies that a KP routine has started or restarted another KPB. Except in the simplest cases, code flow rapidly becomes exceedingly complex as the number of active KPBs increases. In such cases, a work queue model with a dispatcher often removes the need for multiple active KPBs.
- Allocation and deallocation of stacks is sufficiently expensive in system time and resources that consideration should be given to some level of caching KPBs within the application instead.
- Applications using KP services in process context in kernel mode must lock the stacks into the working set.

---

## Glossary

### ABI

Application binary interface. A general term encompassing all of the rules and conventions that programs must follow to run under a particular operating system. For example, the OpenVMS I64 ABI includes the calling standard, register usage, interrupt handling, and other related topics.

According to Intel's *Itanium Software Conventions and Run-Time Architecture Guide*, "an ABI is composed of an API, system-specific conventions, a hardware description, and a run-time architecture."

### ACPI

Advanced configuration and power interface. ACPI provides interfaces for power state management, system device identification, and hardware configuration.

- ACPI specification:

<http://www.acpi.info/spec.htm>

- Intel's ACPI web site:

<http://developer.intel.com/technology/iapc/acpi/>

### ALAT

Advanced load address table. A lookup table on Itanium processor that tracks **speculative** memory accesses and allows the CPU to handle conflicts.

### AML

ACPI machine language. Pseudocode for a virtual machine supported by an ACPI-compatible operating system and in which ACPI control methods and objects are written. See also **ACPI**.

### application registers

128 special-purpose registers used for various functions, labeled ar0 to ar127. For example, ar17 is the register-backing store pointer.

### Aries

HP emulator/translator for HP Precision Architecture (**PA**) application code. Unlike **VEST** (the VAX-to-Alpha translator), Aries emulates only at run-time and does not produce a translated image. When the image exits, all the translated code is discarded and must be regenerated each time the program runs.

### ASL

**ACPI** source language. The programming language for **AML**. ASL source code is compiled into AML images.

**ASM**

Address space match, part of virtual memory management.

**branch registers**

Itanium 64-bit registers, br0 to br7, used to specify the target addresses for indirect branches. The branch registers streamline call/return branching.

**bundle**

The basic instruction format for the Itanium architecture. A bundle is a 128-bit structure consisting of three 41-bit instructions, plus a 5-bit **template**.

**CISC**

Complex instruction set computing. Includes numerous variable-length machine instructions that are more complex than RISC. See also **RISC** and **VLIW**.

**COW**

Copy-on-write. Refers to memory that is mapped by multiple processes provided each process does not attempt to write to it. If a process writes to a page, the page is copied to a process-private location and remapped, leaving the original shared page intact. COW is commonly used by the loader for memory-mapped shared libraries.

**CVF**

Compaq Visual Fortran.

**DSDT**

Differentiated system description table. Part of the **ACPI** subsystem. The DSDT contains the Differentiated Definition Block, which supplies the implementation and configuration information about the base system.

**DVD**

Digital versatile disk.

**DWARF**

Debugging and traceback information (embedded in **ELF**).

**EFI**

Extensible firmware interface. EFI's purpose is to initialize the hardware and start booting an operating system.

For more information, see Intel's EFI Page:

<http://developer.intel.com/technology/efi/index.htm>

**ELF**

Executable and linkable format (ELF) defines the file format for object files, shared libraries, and executable images on Itanium processors. ELF is not specific to the Itanium architecture; it has been used for years on various UNIX systems.

For more information, see the ELF specification from SCO:

<http://stage.caldera.com/developer/gabi/>

## **EPIC**

Explicitly parallel instruction computing. A term coined by Intel to describe the Itanium architecture (comparable to **CISC** or **RISC**). EPIC provides architecturally visible access to parallelism in the CPU. For example, the Intel Itanium processor family defines data speculation and register renaming as part of the architecture, rather than hiding it in the chip implementation as was done with Alpha. EPIC is similar to the very long instruction word (**VLIW**) computing architecture.

## **FADT**

Fixed ACPI description table. A table that contains the **ACPI** hardware register block implementation and configuration details, as well as the physical address of the **DSDT**.

## **FAT**

File allocation table. The disk volume structure used by various Microsoft operating systems and by the Itanium **EFI** operating system loader. There are various FAT variants, including FAT8, FAT12, FAT16, and FAT32. These differ in total permitted volume size and various other changes.

## **fills**

See **spills and fills**.

## **FIT**

Firmware interface table.

## **GEM**

The backend code generator used by most OpenVMS Alpha and I64 compilers.

## **general registers**

Itanium processors have 128 64-bit general registers, numbered as r0 to r127 (sometimes called gr0 to gr127). However, only 32 registers are normally used directly; the other 96 are used for argument passing and to store local variables for the currently executing procedure. Registers r0 to r31 are static, but registers r32 and up are generally managed by the register stack engine (**RSE**) and might be renamed across procedure calls.

## **global pointer (GP)**

The base address of the current global data segment. The base address allows the use of compact, GP-relative addresses in machine code. Because the largest immediate value in an Itanium instruction is 22 bits wide, the global data segment is 4 MB in length. General register 1 holds the current GP.

## **HP OpenVMS Migration Software for Alpha to Integrity Servers**

A tool that translates Alpha executable images to I64 executable images.

## **HWPCB**

Hardware process control block. The portion of the process context that is stored and retrieved by the hardware when a process is rescheduled.

## **HWRPB**

Hardware restart parameter block. These data structures are used to communicate system-level configuration information between the console and OpenVMS both on system startups and system restarts.

## **IA**

Intel architecture.

## **IA-32**

An Intel 32-bit CISC microprocessor architecture. Examples of implementations include the Pentium®, Xeon™, and 80386 series.

## **IA-64**

An old name for the Itanium processor family architecture. The term IA-64 is no longer officially used. However, it still appears in older documentation, and IA64 still occurs in some code.

## **Integrity**

The brand name of the HP server product line based on the Itanium architecture. For example, the rx2600 server is sold as the HP Integrity rx2600.

## **ILP**

Instruction level parallelism. The ability to execute multiple instructions in parallel during the same cycle.

## **instruction group**

A sequence of one or more instructions, delimited by explicit **stops** or taken branches. Instructions within an instruction group must have no **RAW** or **WAW** dependencies. Instruction groups can cross **bundle** boundaries.

## **Intel® Itanium®**

The name of both the 64-bit architecture (the itanium processor family) and the first chip to implement the architecture (the Itanium processor).

## **Intel® Itanium® 2**

The second generation of the Itanium processor family. The **rx2600** is an Itanium 2 system. For more information, refer to the following web site:

[http://www.cpus.hp.com/technical\\_references/ia64.shtml](http://www.cpus.hp.com/technical_references/ia64.shtml)

## **IPINT**

Inter-processor interrupt. Also simply called IPI.

## **IPF**

Itanium processor family.

## **IPMI**

Intelligent platform management interface, a set of APIs used to manage servers. Functions include monitoring environmental data such as temperature, voltage, fans, and chassis switches, as well as alerting, remote console and control, and asset tracking. See the Intel web site for more details:

<http://www.intel.com/design/servers/ipmi/>

**IVT**

Interrupt vector table.

**Jacket routines**

Interface routines that transform the arguments from one calling standard to another.

**MAS**

Multiple address space. In an MAS operating system, all processes on the system have at least a partially private virtual address space. OpenVMS and Tru64 UNIX are both MAS operating systems. See also **SAS**.

**MBR**

Master boot record. The contents of the first disk block and the core data structure of a FAT32 partitionable disk.

**NaT**

Not a thing. An extra bit associated with various registers that indicates whether or not the register contents are valid. Used for **speculative** execution. For example, a read from a nonexistent virtual address would result in NaT. The NaT bit propagates as execution progresses so that any operation involving an NaT results in an NaT.

**NMI**

Non-maskable interrupt.

**NUE**

The Linux Native User Environment, a set of development software used in conjunction with the **Ski** Itanium emulator.

**PA**

Hewlett-Packard (HP) Precision Architecture, a computer architecture also known as PA-RISC. HP is porting its UNIX from PA the to Itanium architecture.

**PAL (Alpha)**

Privileged architecture library. PAL is part of the Alpha architecture. It is a software mechanism for performing low-level operations such as interrupt handling, **TLB** management, and atomic operations that were implemented on VAX as microcode. On Itanium systems, the same functionality is provided as part of the operating system. Whenever possible, the BLISS, C, and Macro compilers for OpenVMS I64 convert CALL\_PAL macros to the equivalent operating system calls for backward compatibility. Not all Alpha PAL operations are implemented on I64; in some cases, programs that call PALcode directly might need to change.

**PAL (Itanium)**

Processor abstraction layer. Part of the Itanium architecture that is implemented in firmware. Provides a consistent interface for processor-specific functions such as hardware errors or system initialization. Generally speaking, **SAL** isolates the operating system from platform-specific implementation differences, while PAL isolates the operating system (and SAL) from processor-specific differences.

**PIC**

Position-independent code. Machine code that uses only relative address references. This strategy allows the code to be loaded to any memory location without modification.

**POSSE**

Pre-operating system startup environment. POSSE is the standard firmware user interface on **Itanium® 2** servers. It is an HP value-added component that is available only on HP systems.

**PPN**

Physical page number.

**PTE**

Page table entry.

**predication**

The conditional execution of instructions based on a set of special 1-bit registers called predicates. Most I64 instructions include a predicate number. When the corresponding predicate register is true (1), the instruction is executed. When it is false (0), the instruction is treated as a “no-op.”

**RAW**

Read-after-write violation. A type of data dependency between two instructions in one instruction group. The later instruction reads data from the same register to which the earlier instruction wrote.

**RID**

Region identifier.

**RISC**

Reduced instruction set computing. Fixed-size machine instructions that are intended to be simpler machine instructions than CISC. See also **CISC** and **VLIW**.

**RSE**

Register stack engine. An on-chip entity that handles register renaming and **spills and fills** across procedure calls.

**rx2600**

An HP server with up to two Itanium 2 CPUs. The rx2600 is the first platform for customer use to run the OpenVMS Itanium port.

**SAL**

System abstraction layer. Part of the Itanium architecture, implemented in firmware. Provides a consistent interface for platform-specific functions such as hardware errors or system initialization. Generally speaking, SAL isolates the operating system from *platform*-specific implementation differences, while **PAL** isolates SAL and the operating system from *processor*-specific differences.



## **SAS**

Single address space. In an SAS operating system, all processes share the same address space. See also **MAS**.

## **SCI**

System control interrupt. A system interrupt used by hardware to notify the operating system about **ACPI** events.

## **Ski**

A free Itanium emulator for Linux. Ski emulates the Intel Itanium Processor Family architecture but not any particular implementation.

## **speculation**

The process of moving memory accesses earlier in the code to hide memory access latencies. If it turns out that the memory access was invalid or unnecessary, it can easily be discarded. The Itanium processor instructions `ld.a`, `ld.s`, `chk.a`, and `chk.s` are used for speculation.

## **spills and fills**

The rotating portion of the Itanium register file that is used as a stack. If it overflows, registers are saved to a memory region called the register backing store. A *spill* flushes one or more registers out to the backing store; a *fill* restores registers from the backing store. Spills and fills are handled automatically by the register stack engine (**RSE**) on the CPU without program intervention.

## **SRM console**

Firmware for Alpha systems that includes boot support, PALcode, and platform-specific functionality used by OpenVMS, UNIX, and Linux. The chevron prompt (`>>>`) that you see when you power on any Alpha system is part of the SRM console. The console behavior is defined in the *Alpha System Reference Manual* (SRM).

The Itanium architecture does not have an equivalent to the SRM console. Instead, the SRM functionality is split between the **PAL/SAL** firmware, the **EFI** boot software, and the operating system.

## **stop**

The end of an **instruction group**, designated in by a double semicolon (`::`) in assembly listings.

## **Superdome**

The high-end HP server product line. Superdome servers support about 64 CPUs, either PA-RISC or Itanium. Eventually, all new Superdome systems will be Itanium based.

## **SWIS**

SoftWare interrupt support. The set of services that implement the interrupt model used by OpenVMS. This includes processor IPL, ASTs, mode changes, and software interrupts. On Alpha, this support was implemented in PALcode.

**template**

On Itanium processors, templates define the combinations of instructions that can be placed into a single **bundle**. For example, a bundle with the MFI template must have, in this order, a memory instruction (M), a floating-point instruction (F), and an integer instruction (I). Templates also define the location of **stops**. 24 templates are defined in the Itanium architecture, but there are only 12 combinations of letters. For each combination, there are two template versions: one with a stop at the end, and one without a stop. Some templates also contain stops within the bundle.

**TLB**

Translation lookaside buffer. An on-chip cache that stores recently used mappings between real and virtual memory. The Itanium architecture has separate TLBs for instructions and data.

**USB**

Universal Serial Bus. A connector for I/O devices such as keyboards, mice, printers, cameras, external disks, and so forth.

**VEST**

A tool that translates VAX executable images to Alpha executable images. Also known as DECMigrate.

**VHPT**

Virtual hash page table.

**VIAL**

OpenVMS to itanium abstraction layer. VIAL is the equivalent of the OpenVMS SRM console PALcode layer.

**VLIW**

Very long instruction word computing architecture. Exceedingly complex instructions reminiscent of the microcode engine found deep within many VAX processors. See also **CISC**, **RISC**, and **EPIC**.

**VPN**

Virtual page number.

**VRN**

Virtual region number.

**WAW**

Write-after-write violation. A type of data dependency between two instructions in one instruction group. Both of the instructions write to the same destination register.

**WWID**

Worldwide ID. A globally unique Fibre Channel device identifier analogous to an Ethernet station address.

## A

---

- Ada, 6–1, 6–2
- ADA 83, A–4
- ADA compiler
  - replacement, B–1
- Alpha instructions
  - interpreting, 4–4
- Analyze/Image utility (ANALYZE/IMAGE), 5–2
- Analyze/Object utility (ANALYZE/OBJECT), 5–2
- Application
  - size, A–3
- Applications
  - and compilers, A–4
  - architecture dependencies, A–6
  - assessing, 3–2
    - contents being ported, 3–4
    - dependencies, 3–5, A–4
    - operational tasks, 3–6
    - porting method, 3–2
  - calling standard, 3–4, A–8
  - components, A–3
  - development environment, 3–6
  - documentation, A–4
  - establishing baseline values for, 4–7
  - hardware dependencies, A–5
  - languages used, A–4
  - machine instructions, A–8
  - operating environment, 3–6
  - performance sensitivity, A–8
  - requiring changes, 3–4
  - software dependencies, 3–5, A–5
  - synchronization among parts of, A–7
  - use of global sections, A–7
- Architecture dependencies, A–6
- asm statement, A–9
- ASTs (asynchronous system traps), A–7

## B

---

- Baseline values for application
  - establishing, 4–7
- BASEstar Classic
  - replacement, B–1
- BASIC, 2–5, 6–1, 6–2

- Binary translator, 4–1, 5–3
- BLISS, 6–1
- BLISS builtin, A–9
- BLISS compilers
  - differences between Alpha and I64, 6–2
- Bugs
  - latent, 4–9

## C

---

- C, 2–5, 6–1, 6–18
  - include files, 4–3
- C++, 2–5, 5–10, 6–1, 6–19
- Calling standard, 3–4, A–8
  - changes for I64, 2–1
  - code that relies on, 4–18
- C asm statement, A–9
- CLI\$DCL\_PARSE
  - external definition, 4–15
- CLUE (Crash Log Utility Extractor)
  - See Crash Log Utility Extractor
- CMA\_DELAY API library routine, 4–17
- CMA\_TIME\_GET\_EXPIRATION API library routine, 4–17
- CMS, 4–3
- COBOL, 2–5, 5–1, 5–10, 6–1, 6–15
- Code
  - assembler, A–8
  - requiring changes, 3–4
  - user-written, A–8
- Code Management System
  - See CMS
- command definition file, 4–15
- Compile commands
  - changes required, 4–7
- Compile procedures, 4–3
- Compilers
  - assessing porting support, A–4
  - commands, 4–7
  - native Alpha, 5–1
  - native I64, 5–1
  - optimizing, 5–1
  - preparing to port applications, 6–1
  - qualifiers for VAX floating point data types, 5–1

- Conditionalized code, 3-4, 4-10, A-8
- Condition handlers, 4-13, A-7
- Connect-to-interrupt mechanisms, A-7
- Crashes
  - analyzing, 4-7
- Crash Log Utility Extractor (CLUE), 4-8, 5-2
- \$CRMPSC, A-7

## D

---

- Data alignment, 3-5, 4-17 to 4-18, A-6
  - compiler options, 4-18
  - finding unaligned data, 4-18
  - incompatibility with translated software, 4-18
  - performance, 4-17
- Data structures
  - alignment, 3-5, A-6
  - floating-point, 3-4, 3-7, A-6
  - internal, 3-4, A-7
  - mechanism array, 3-4, A-7
- Data types
  - floating-point, 3-4, 3-7, A-6
- Debuggers, 4-7
  - native I64, 5-2
  - XDelta, 5-2
- Debugging, 5-2
- Debug symbol table, 4-13
  - format, 3-4, A-9
- DECset, 5-2
- Digital Test Manager, A-5
- Drivers
  - terminal, 3-5
- Dump files
  - See System dump files
- DWARF image file format, 4-13

## E

---

- ELF object file format, 4-13
- Executable image file format, 3-4, A-9

## F

---

- Floating-point data types, 3-4, 3-7, A-6
  - IEEE, 4-13
  - VAX, 4-13, 5-1
- Fortran, 6-1, 6-15
- Fortran 77, A-4

## H

---

- HP OpenVMS Migration Software for Alpha to Integrity Servers, 4-3, 5-3
  - analytical ability, 4-5
  - capabilities, 4-4
  - resources required, 4-3
  - runs on Alpha and I64 systems, 4-3

- HP rx2600 system, 4-2
- HP Services support, 3-8

## I

---

- Include files
  - for C programs, 4-3
- Internal data structures, 3-4, A-7
- Interoperability
  - confirming, 4-9

## J

---

- Jacket routines, 4-4, 4-5
- Java, 6-1, 6-20
- Jump to subroutine (JSB), 3-5

## K

---

- KPB, C-3
- KP control routines, C-9
- KP services, C-1

## L

---

- Languages, A-4
- Layered products
  - unsupported, B-1
- LIB\$LOCK\_IMAGE routine, 4-19
- LIB\$UNLOCK\_IMAGE routine, 4-19
- LIB\$WAIT
  - common code for I64 and Alpha, 4-14
- Librarian utility (LIBRARIAN)
  - native Alpha, 5-2
- Library interfaces, 3-1, 3-4
- Linkage declarations, A-9
- Link commands
  - changes required, 4-7
- Linker utility
  - native I64, 5-2
- Link procedures, 4-3
- Logical names
  - for tools and files, 4-3

## M

---

- Machine instructions
  - correct execution of, A-8
- Macro-32, 5-1, 5-10, 6-1
- Macro-32 compiler, 5-1
- Macro-32 compiler, 6-20
- Macro-64, A-4
- Macro-64, 6-1
- MACRO code
  - replacing, A-4
- Mechanism array data structure, 3-4, 4-13

Message utility (MESSAGE)  
    native Alpha, 5-2  
Migration environment  
    setting up, 4-2  
Migration tools, 4-3  
MMS, 4-3  
Module Management System  
    See MMS  
Multiprocessing, A-7

## O

---

Object file format, 3-4, A-9  
    reliance on, 4-13  
OTS\$CALL\_PROC routine, 4-5

## P

---

Pascal, 2-5, 6-1, 6-21  
Pathworks 32  
    replacement, B-1  
PGFIPLHI bugchecks, 4-19  
PL/I, A-4  
Porting applications, 6-1  
Porting checklist, 7-4  
Porting dependencies  
    external, 3-5, A-4  
Porting method, 3-2  
Porting needs  
    assessing, 3-1  
Porting resources, 3-8  
Porting steps, A-9  
Privileged mode operation, A-7  
Programming languages  
    Ada, 6-1, 6-2  
    BASIC, 5-1, 6-1, 6-2  
    BLISS, 5-1, 6-1, 6-2  
    C, 4-3, 5-1, 6-1, 6-18, 6-19  
    C++, 5-1, 6-1  
    COBOL, 5-1, 6-1, 6-15  
    Fortran, 5-1, 6-1, 6-15  
    Java, 6-1, 6-20  
    Macro-32, 6-1  
    Macro-64, 6-1  
    Pascal, 5-1, 6-21  
    VAX MACRO, 5-1, 6-20  
Program sections  
    overlaid, 4-12

## R

---

Recompiling  
    changes in compile commands, 4-7  
    resolving errors, 4-7  
Relinking, 5-2  
    changes in link commands, 4-7

RSE stack, C-1

## S

---

SDA  
    See System Dump Analyzer utility  
SS\$\_FLTDIV exception, 4-13  
SS\$\_FLTINV exception, 4-13  
SS\$\_HPARITH exception, 4-13  
Stacks  
    switching, C-1  
Synchronization  
    of interprocess communication, A-6  
SYS\$BASE\_IMAGE  
    linking against, A-7  
SYS\$GOTO\_UNWIND system service, 4-12, A-6  
SYS\$GOTO\_UNWIND\_64 system service, 4-12  
SYS\$LCKPAG system service, 4-20  
SYS\$LCKPAG\_64 system service, 4-20  
SYS\$LKWSET system service, 4-19, A-6  
SYS\$LKWSET\_64 system service, 4-19  
System-Code Debugger, 4-7  
    See also Debugger  
System Dump Analyzer utility (SDA), 5-2  
    OpenVMS Alpha, 4-8  
System dump files  
    analyzing, 4-8  
System services, 3-1, 3-4  
    SYS\$GOTO\_UNWIND, 4-12, A-6  
    SYS\$GOTO\_UNWIND\_64, 4-12  
    SYS\$LCKPAG, 4-20  
    SYS\$LCKPAG\_64, 4-20  
    SYS\$LKWSET, 4-19, A-6  
    SYS\$LKWSET\_64, 4-19

## T

---

Terminal drivers, 3-5  
Testing applications  
    compile, A-9  
    establishing baseline values on Alpha systems,  
        4-7  
    on I64 systems, 4-8  
    regression testing, A-10  
    regression tests, A-5  
    system testing, A-10  
    unit testing, A-10  
Test tools, 4-3  
    I64 specific, 4-9  
    ported to I64, 4-8  
THREADCP command, 4-16  
Threading-like code, 3-5  
Thread interfaces  
    legacy API library routines, 4-17  
    support on I64, 4-16  
TIE (Translated Image Environment), 4-3  
    description, 4-4  
    invoked automatically, 4-4

Translated images  
    coexistence, 4–5  
    contents, 4–5  
Translating, 3–2, 5–3, A–3, A–4  
    programs in languages with no I64 compiler,  
        5–3  
    tools for, 4–4  
    type of image produced, 4–5  
Translator

See Binary translator

---

## V

VAX MACRO  
    recompiling on OpenVMS I64 systems, 5–1

---

## W

Working set  
    modifying, A–6

---

## X

XDelta Debugger, 4–7, 5–2